

StorageTek Enterprise Library Software

Configuring HSC and VTCS

Release 7.3 for MVS

E60318-06

May 2019

StorageTek Enterprise Library Software Configuring HSC and VTCS, Release 7.3 for MVS

E60318-06

Copyright © 2015, 2019, Oracle and/or its affiliates. All rights reserved.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information about content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services unless otherwise set forth in an applicable agreement between you and Oracle. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services, except as set forth in an applicable agreement between you and Oracle.

Contents

Preface	xxi
Audience	xxi
Prerequisites	xxi
Documentation Accessibility	xxi
What's New?	xxiii
1 Planning for HSC/VTCS Configuration	
How VSM Measures Sizes and Capacities	1-2
Determining HSC/VTCS Configuration Values	1-3
Control Data Sets	1-3
Control Data Set Placement	1-4
Control Data Set Sharing	1-4
Sharing Requirements	1-4
Sharing Recommendations	1-4
Additional Recommendations When HSC Runs on an MVS Guest Under VM	1-5
Additional Recommendations When the CDS is Shared Between MVS and VM	1-6
Limiting CDS Access Privileges	1-6
Serializing CDSs	1-6
Set Utility Serialization	1-7
GRS Serialization	1-7
Multiple HSC TapePlex Considerations	1-8
Example	1-8
HSC CDS Performance and Sharing Tips	1-8
Unicenter CA-MIM/MII Considerations	1-9
GRS Considerations	1-9
When HSC RESERVEs Must Remain as RESERVEs (All Environments)	1-9
When You May Want to Leave HSC RESERVEs as RESERVEs	1-10
When HSC RESERVEs Must be Converted to Global Enqueues	1-11
How to Leave HSC RESERVEs as RESERVEs in a GRS Environment	1-11
How to Leave HSC RESERVEs as RESERVEs in an MIM/MII Environment	1-12
How to Convert HSC RESERVEs to Global Enqueues in a GRS Environment	1-13
How to Convert HSC RESERVEs to Global Enqueues in an MIM/MII Environment	1-13
CDS DASD Space Requirements	1-14
Calculating DASD Space Requirements for POOLPARM/VOLPARM	1-15

Calculating DASD Space Requirements for VTCS	1-15
Calculating DASD Space Requirements for VAULTVOL	1-15
Calculating DASD Space Requirements for CDKLOCK.....	1-16
CDS VTCS Level Requirements	1-16
Specifying Valid CDS Levels for VTCS 7.3	1-17
Backing Out From an I Level CDS.....	1-18
Using CDS Logging to Help Recover the CDS	1-18
Planning for LSMs.....	1-19
Placement of ACS Network Connections on Your Network	1-19
LSM Unit Address Requirements.....	1-19
LSM Pass-Thru Port (PTP) Relationships.....	1-19
Host-to-Host Communication Options	1-20
Calculating Cartridge Capacity - SL8500 and SL3000	1-20
Planning for the SL8500	1-20
Planning for Library-Attached Transports.....	1-21
Planning for Library-attached Transports as RTDs	1-22
Planning for Library Volumes.....	1-23
Guidelines for MVCs	1-23
VTCS Considerations to Correctly Specify MVC Media.....	1-23
Examples	1-26
Using the STORclas MEDIA Parameter for MVC Media Preferencing.....	1-26
Planning for VTSSs	1-28
Planning for VTDs	1-29
Planning for VTVs.....	1-30
Data Chaining a VTD Read Forward or Write Command.....	1-31

2 Planning VTCS Operating Policies

Making CDS Logging Required.....	2-1
Usage Notes	2-1
Running IMPORT/EXPORT When Logging Required is Enabled	2-2
Locked Resource Policy.....	2-2
VTSS Policies	2-2
AMT Settings	2-3
Usage Notes	2-3
Deleting Scratched VTVs	2-4
Usage Notes	2-4
VTV Page Size.....	2-5
Usage Notes	2-5
Maximum VTV Size.....	2-6
Usage Notes	2-6
Maximum RTDs per VTSS.....	2-7
Usage Notes	2-7
RTD/CLINK Addresses - Maximum 32 RTDs.....	2-7
General VTV Policies	2-10
Maximum VTVs per MVC.....	2-11
Usage Notes	2-11

Hosts Disabled from Migration, Consolidation, and Export by VTV or Management Class.....	2-11
Usage Notes	2-11
Recall VTVs with Read Data Checks.....	2-12
Usage Notes	2-12
Early Time to First Byte (ETTFB)	2-12
Usage Notes	2-12
VTV VOLSAFE Policy	2-13
Usage Notes	2-13
VTV Migration Policies	2-14
Streamed/Stacked Migrates	2-14
Usage Notes	2-15
Minimum Residency Time before Automatic Migration Candidacy.....	2-16
Usage Notes	2-16
Maximum Residency Interval before Deletion	2-17
Usage Notes	2-17
Immediate Migration Delay Interval	2-17
Usage Notes	2-17
Maximum and Minimum Concurrent Migration Tasks	2-18
Usage Notes	2-18
Prioritizing Automatic and Immediate Migration Workloads	2-19
Usage Notes	2-19
Controlling Migration Workloads	2-19
Usage Notes	2-20
Conditionally Change Migration Preferencing/Depreferencing	2-20
Usage Notes	2-20
Dependent Immediate Migration	2-21
Usage Notes	2-21
Putting It All Together: VTV Migration Control Example for an Enterprise.....	2-21
Enabling the Remote and Local ACS Migrations.....	2-22
Migration Policies for a Vault Solution.....	2-23
Enabling the "Other Data" Migrations	2-24
VTV Replication Policies	2-25
VTCS Replication - Always or Changed VTV Only.....	2-25
Usage Notes	2-25
VTCS Replication - Synchronous or Asynchronous	2-26
Usage Notes	2-26
VTCS Enhanced Replication.....	2-27
Usage Notes	2-27
VTD Mount Failures.....	2-28
VTD Dismount/VTV Audit Failures	2-28
Vary CLInk Processing.....	2-28
MVC Policies	2-29
Standard Reclaim Threshold - Determines MVC Eligibility for Reclamation (Requires Recall/Remigration) 2-30	
Usage Notes	2-30
Dynamic Reclaim Threshold - No MVC Mount or VTV Recall/Remigration Occurs	2-32

What is Dynamic Reclaim?	2-32
Usage Notes	2-33
Dynamic Reclaim - Specifying VTV Data Movement.....	2-33
Usage Notes	2-34
VLE Reclaim Threshold - Determines VMVC Eligibility for Reclamation.....	2-38
Usage Notes	2-38
Free MVCs Threshold - Starts Automatic Space Reclamation	2-38
Usage Notes	2-39
Eligible/Total MVCs Threshold - Starts Automatic Space Reclamation.....	2-39
Usage Notes	2-39
Maximum MVCs Processed Per Reclaim	2-40
Usage Notes	2-40
Maximum MVCs Concurrently Processed for Reclamation and Drain.....	2-41
Hosts Disabled from Reclamation	2-41
Usage Notes	2-41
MVC Retain Interval.....	2-41
Usage Notes	2-41
MVC Mount Timeout Interval	2-42
Non-Library Resident MVC Policies.....	2-42
MVC Initialization on First Mount	2-43

3 Preparing for Configuration

Defining A Security System User ID for HSC and VTCS	3-1
Configuring MVS Device Numbers and Esoterics	3-1
Assigning MVS Device Numbers to VTDs and Library Drives.....	3-2
Assigning MVS Device Numbers to RTDs.....	3-2
Setting the MVS Missing Interrupt Handler (MIH) Value	3-2

4 Configuring HSC

Creating the HSC CDS	4-1
Coding the LIBGEN Macros.....	4-1
Running the SLICREAT Utility To Format the New CDS	4-6
Running the SLUADMIN CDSCREAT Utility to Create the CDS (Tapeless VSM Only)	4-8
Defining and Creating CDS Log Files.....	4-9
Running Multiple HSC Tapeplexes on One LPAR (MULT Mode)	4-10
Defining Volumes to ELS	4-10
Why Should You Use POOLPARAM/VOLPARAM?.....	4-11
Some Things to Consider with POOLPARAM/VOLPARAM.....	4-12
Defining Library Volumes	4-13
Defining Cleaning Cartridges	4-13
Defining MVCs.....	4-14
Defining VTVs	4-14
Validating and Applying the Volume Definitions	4-15
Adding Definitions for ACF/VTAM Communications	4-15
Defining the SYS1.PARMLIB Member SMFPRMxx	4-16
Creating the HSC PARMLIB Member	4-17
EXEC Parm Statement.....	4-18

Syntax	4-18
Parameters.....	4-18
Defining Command Authority	4-19
Updating HSM.....	4-19
Creating and Cataloging the HSC Startup Procedure.....	4-19
Running Multiple HSC Tapeplexes on One LPAR (MULT Mode)	4-20
EXEC Statement	4-21
Syntax	4-21
Parameters	4-21
Sample HSC Startup Procedure.....	4-23
DD Statements for the HSC Startup Procedure.....	4-24
Specifying TCPIP Options.....	4-24
Performance Considerations	4-24
Define Dispatching Priority for ELS Software	4-24
Reduce Pass-Thrus	4-25
Unavoidable Pass-Thrus	4-25
Unnecessary Pass-Thrus	4-25
Scheduled Pass-Thrus	4-25
Methods to Reduce Pass-Thru Activity	4-25
Reduce Operator Intervention	4-26
Reduce Scheduling Contention	4-28

5 Reconfiguring a TapePlex

When Do You Need to Reconfigure Your TapePlex?.....	5-1
The MERGEcdfs Utility.....	5-2
Basic MERGEcdfs Procedure	5-2
Create a New, Empty CDS	5-2
Prepare to Run MERGEcdfs	5-4
Run the MERGEcdfs Utility.....	5-5
Ensure the Merge was Successful.....	5-6
Consolidating Two Data Centers.....	5-6
Combining Two Separate ACSs into One ACS	5-7
Combining Two Separate Streamline Libraries.....	5-8
Configurations for ACSs Numbered Right to Left	5-8
Configurations for ACSs Numbered Left to Right	5-9
SLSMERGE Statements - ACSs Numbered Right to Left	5-9
SLSMERGE Statements - ACSs Numbered Left to Right	5-10
Removing PTPs and Splitting an ACS Using MERGEcdfs	5-11
Renaming a VTSS.....	5-11
Removing Vaulted Volume Information.....	5-12
Reconfiguring a Tapeplex Without MERGEcdfs	5-12
Using Dynamic Reconfiguration	5-12
Adding LSMs.....	5-13
Dynamically Reconfiguring Streamline Library Transports	5-13
Adding Transports.....	5-14
Removing Transports	5-15
Replacing Transports.....	5-15

Changing Transport Unit Addresses	5-16
Adding SL8500 or SL3000 ACSs	5-16
Deleting an SL8500 or SL3000 ACS	5-17
Adding and Removing SL8500 or SL3000 Expansion Modules	5-17
Changing SL8500 CAP Types - HSC Down	5-18
Changing SL8500 CAP Types - HSC Up	5-19
Using Audit Instead of MERGEcds for Streamline Libraries	5-19
Adding a New VTSS	5-20
HSC SET VAULTVOL and SET VAULT Utilities	5-20
HSC SET CDKLOCK Utility	5-20
Partitioning Streamline Libraries	5-20
Partitioning Requirements and Prerequisites	5-21
Restrictions	5-21
CAP Considerations	5-22
LibraryStation Considerations	5-22
Definitions	5-23
Host Group	5-23
SL8500 Enhanced Partition	5-23
SL3000 Partition	5-24
Partitioning Procedures	5-26
Starting from Scratch – LIBGEN, SLICREAT	5-26
Converting a Non-partitioned Library to a Partitioned Library	5-28
Converting a Partitioned Library to a Non-partitioned Library	5-30
Adding a Library to the HSC Complex	5-33
Removing a Library from the HSC Complex	5-36
Adding a Partition to the Library	5-38
Removing a Partition from the Library	5-41
Moving a Partition from One Host Group to another Host Group	5-44
Adding Resources to a Library	5-47
Removing Resources from a Library	5-51
Changing the CAP State for Bulk CAP - HSC Down	5-54
Changing the CAP State for Bulk CAP - HSC Up	5-55
Varying the ACS Online After a Configuration Change	5-55

6 Configuring VTCS

Building a Simple CONFIG Deck	6-2
To Build a Simple CONFIG Deck:	6-2
Updating the Tape Management System for VSM	6-4
Defining MVC Pool Volser Authority	6-5
Storing VTCS Locks in a Coupling Facility (Optional)	6-6
When to Implement VTCS Locks	6-6
Requirements	6-7
Sizing the Coupling Facility Structure	6-7
Defining the Coupling Facility Structure to MVS	6-8
Reverting to Storing VTCS Lock Data in the CDS	6-9
Managing Failures/Unavailability of the VTCS lock structure	6-9
Managing Notification of Resource Contention	6-9

Configuring VTCS for 512 VTD Support	6-9
Enabling 512 VTD Support.....	6-10
Updating or Disabling 512 VTD Support.....	6-10

7 Starting and Stopping HSC

Starting HSC	7-1
HSC Startup Scenarios.....	7-2
IPLing the System	7-3
Preinitializing HSC as an MVS Subsystem	7-3
Starting HSC with PARM='INIT'	7-3
Initializing HSC Under the Master Subsystem.....	7-4
Starting HSC As a Subsystem	7-4
Specifying the HSC Service Level at Startup	7-4
Starting the HSC Behavior Change	7-4
Configuration Mismatches	7-5
Stopping HSC	7-5

A HSC LIBGEN Macros

SLIACS	A-1
Syntax.....	A-1
Parameters.....	A-1
SLIALIST	A-3
Syntax.....	A-4
Parameters.....	A-4
SLIDLIST	A-4
Syntax.....	A-4
Parameters.....	A-5
SLIDRIVS	A-5
Syntax.....	A-5
Parameters.....	A-6
SLIENDGN	A-8
Syntax.....	A-9
Parameters.....	A-9
SLILIBRY	A-9
Syntax.....	A-9
Parameters.....	A-9
SLILSM	A-14
Syntax.....	A-15
Parameters.....	A-15
SLIRCVRY	A-20
Syntax.....	A-20
Parameters.....	A-21
SLISTATN	A-22
Syntax.....	A-22
Parameters.....	A-23

B LIBGEN Macro Samples

One Host, Two ACSs, One 9310 LSM, One SL8500 Library Configuration	B-1
Four Host, Two ACSs, Two SL8500 Libraries Configuration	B-4
Three Host, Three ACSs, Two SL8500 Libraries, One SL3000 Library Configuration	B-7
Three Host, Two ACSs, Two SL3000 Libraries Configuration	B-10
One Host, One ACS, One 9310 LSM, Dual LMU Configuration	B-11
Two Host, One ACS, Two 9310 LSM Configuration	B-13

C Special Cases: Advanced Uses of CONFIG

CONFIG Example: All Hosts Access VTDS in One VTSS, Only Selected Hosts Access VTDS in Second VTSS	C-1
CONFIG Example: Update Configuration to Add RTDs	C-2
CONFIG Example: Update Configuration to Add MVCs and VTVs and Change AMTs	C-2
CONFIG Example: Denying Host Access to a Physically Removed VTSS	C-3
CONFIG Examples: Using the Maximum 32 RTDs Feature	C-4
Implementing Support for Maximum of 32 RTDs	C-4

D Configuring SL8500 Library Communications

TCP/IP Communications - Important Considerations	D-1
SL8500 TCP/IP Connections	D-1
Shared Networks	D-1
Dual IP Connections	D-2
Two SL8500 Network Connections - Two Mainframe IP Addresses	D-2
Sample Configuration - Two Dedicated Routes	D-5
Setup Statements and Displays	D-6
HSC LMUPDEF PARMLIB Member	D-6
TCP/IP Profile Data Set	D-6
TCP/IP Console Displays	D-7
HSC Display Showing SL8500 Connection	D-8
Two SL8500 Network Connections - One Mainframe IP Address	D-9
Sample Configuration - One Host IP, Two SL8500 Network Connections	D-11
Setup Statements and Displays	D-12
HSC LMUPDEF PARMLIB Member	D-12
TCP/IP Profile Data Set	D-12
TCP/IP Console Displays	D-13
HSC Display Showing SL8500 Connection	D-14
Multiple SL8500 Library Connections	D-15
Connecting to Multiple SL8500 Libraries in an ACS	D-15
Sample Configuration - Four SL8500 Network Connections, Four Mainframe IP Addresses	D-15
Powering Down the SL8500 - HSC Requirements	D-16
Multiple TCP/IP Redundant Electronics	D-17
Sample Configuration #1 - One SL8500 Network Single Redundant Pair Connections, Four Mainframe IP Addresses	D-17
Sample Configuration #2 - One SL8500 Network Dual Redundant Pair Connections, Four Mainframe IP Addresses	D-19

Sample Configuration #3 - One SL8500 Network Dual Redundant Pair and 2 Dual TCP/IP Connections, Four Mainframe IP Addresses	D-21
Troubleshooting Library Connectivity Problems	D-23
Basic Tracing	D-23
Packet Tracing	D-24
E Configuring SL3000 Library Communications	
TCP/IP Communications - Important Considerations	E-1
Shared Networks	E-1
SL3000 Dual IP Connections	E-2
Two SL3000 Network Connections - Two Mainframe IP Addresses	E-2
Sample Configuration - Two Dedicated Routes	E-5
Setup Statements and Displays	E-6
HSC LMUPDEF PARMLIB Member	E-6
TCP/IP Profile Data Set	E-6
TCP/IP Console Displays	E-7
HSC Display Showing SL3000 Connection	E-8
Two SL3000 Network Connections - One Mainframe IP Address	E-9
Sample Configuration - One Host IP, Two SL3000 Network Connections	E-11
Setup Statements and Displays	E-12
HSC LMUPDEF PARMLIB Member	E-12
TCP/IP Profile Data Set	E-12
TCP/IP Console Displays	E-13
HSC Display Showing SL3000 Connection	E-14
Troubleshooting Library Connectivity Problems	E-15
F VSM4 ESCON Configuration	
VSM4 with 32 Ports	F-2
VSM4 Configuration Examples - 32 Ports	F-5
VSM4 Configuration Example: 16 Host Ports, 16 RTD Ports	F-5
CONFIG Example for VSM4 with 16 Host Ports, 16 RTD Ports	F-6
VSM 4 Configuration Example: 20 Host Ports, 12 RTD Ports	F-6
CONFIG Example for VSM4 with 20 Host Ports, 12 RTD Ports	F-7
VSM4 with 16 Ports	F-7
VSM4 Configuration Examples - 16 Ports	F-9
VSM4 Configuration Example: 8 Host Ports, 8 RTD Ports	F-10
CONFIG Example for VSM4 with 8 Host Ports, 8 RTD Ports	F-10
VSM 4 Configuration Example: 10 Host Ports, 6 RTD Ports	F-10
CONFIG Example for VSM4 with 10 Host Ports, 6 RTD Ports	F-11
IOCP Example for Single MVS Host Connected to a VSM4 Using ESCON Directors	F-12
Logical Paths for VSM 4 with 32 Ports	F-13
G VSM4 FICON Front-End and Back-End Configuration	
VSM4 FICON VCF Card Options	G-1
FICON/ESCON Best Practices	G-5
VSM4 FICON Front-End and Back-End Configuration Examples	G-6

VSM4 Configuration Example: 8 VCF Cards, FICON Directors, 16 RTDs.....	G-6
CONFIG Example for VSM4 FICON with 8 VCF Cards, FICON Directors, 16 RTDs.....	G-7
VSM4 Configuration Example: 8 VCF Cards, 4 CLINKs, FICON Directors for 8 RTDs	G-8
CONFIG Example for Bi-Directional Clustered VSM4 FICON Back-End	G-8
IOCP Example for Single MVS Host Connected to a VSM4 Using FICON Directors.....	G-10

H VSM5 FICON Configuration

VSM5 FICON VCF Card Options - Maximum 16 RTDs	H-1
VSM5 FICON VCF Card Options - Maximum 32 RTDs	H-3
FICON Port Processing	H-4
FICON Port Operations Best Practices.....	H-5
VSM5 FICON Front-End and Back-End Configuration Examples.....	H-6
Implementing Support for Maximum of 32 RTDs	H-6
VSM5 Configuration Example: 8 VCF Cards, FICON Directors, 32 RTDs.....	H-7
CONFIG Example for VSM5 FICON with 8 VCF Cards, FICON Directors, 32 RTDs.....	H-8
VSM5 DOP Panels for Maximum 32 RTDs	H-9
Channel Configuration Status Screen	H-9
Channel Configuration and RTD Path Validation Screen	H-10
Real Tape Drive Status Screen	H-12
IOCP Example for Single MVS Host Connected to a VSM5 Using FICON Directors	H-12

I VSM5 ESCON/FICON Configurations

VSM5 ICE/VCF Card Options	I-1
VSM5 Configuration Example: 8 ICE Cards, 16 Host Ports, 16 RTD Ports.....	I-5
CONFIG Example for VSM5 with 16 Host Ports, 16 RTD Ports	I-6
IOCP Example for Single MVS Host Connected to a VSM5 Using ESCON Directors.....	I-7

J VSM5 New Models

K VSM 6/VSM 7 Configuration

Reconfiguring VTCS	K-1
Configuration Examples for Single MVS Host Connected to a VSM 6/VSM 7 Using FICON Directors	K-5
IOCP Example	K-6
Using HCD to Define a VSM 6/VSM 7	K-8
Define the FICON Channels.....	K-9
Define the VSM 6/VSM 7 As Sixteen 3490 Images and Connect Them to the FICON Channels	K-9
Define 16 VTDs for Each Control Unit	K-12

L Tapeless VSM

How Does Tapeless VSM Work?	L-1
Tapeless VSM Example.....	L-2
Configuring the System.....	L-2
Defining Policies.....	L-3
Using CTR in "Tapeless" and "With Tape" Environments	L-3

LIBGEN Example for Tapeless ACS..... L-4

M Running the oVTCS CDS Database Server

Glossary

Index

List of Figures

2-1	VSM5 with 8 VCF cards - Max 32 RTDs	2-8
2-2	Standard Reclaim Example.....	2-32
2-3	Blank Partitioned MVC	2-32
2-4	Full Partitioned MVC	2-33
2-5	Dynamically Reclaimed Partitioned MVC.....	2-33
2-6	Dynamically Reclaimed Partitioned MVC - New Migrations.....	2-33
2-7	Three ALP MVCs and MOVEDATA(None).....	2-34
2-8	MVC Report After MOVEDATA(None) Specified	2-34
2-9	Two ALP MVCs and MOVEDATA(Minimum).....	2-35
2-10	MVC Report After MOVEDATA(Minimum) Specified	2-36
2-11	VTVs Before Reclaim	2-36
2-12	VTVs After Reclaim	2-37
2-13	Three VLE MVCs with MOVEDATA(Minimum).....	2-37
2-14	VLE MVC Report After MOVEDATA(Minimum) Specified	2-38
4-1	EXECPARM Statement	4-18
4-2	EXEC Statement	4-21
5-1	Existing Right to Left Configuration - Three Separate ACSs	5-8
5-2	Desired Configuration - Single ACS	5-9
5-3	Existing Left to Right Configuration - Three Separate ACSs	5-9
5-4	Desired Configuration - Single ACS	5-9
5-5	SLSMERGE Statements Right to Left	5-10
5-6	SLSMERGE Statements Left to Right.....	5-10
5-7	SL8500 Partitioning Example	5-24
5-8	SL3000 Partition and Host Group Example for Cells and CAPs	5-25
5-9	SL3000 Partition and Host Group Example for Drives	5-26
A-1	SLIACS Syntax.....	A-1
A-2	SLIALIST Syntax	A-4
A-3	SLIDLIST Syntax	A-4
A-4	SLIDRIVS Syntax	A-5
A-5	SL3000 Drive Panel Layout.....	A-7
A-6	SLIENDGN Syntax	A-9
A-7	SLILIBRY Syntax	A-9
A-8	SLILSM Syntax	A-15
A-9	Example of Pass-thru Port Relationships	A-18
A-10	SLIRCVRY Syntax.....	A-21
A-11	SLISTATN Syntax	A-22
B-1	One Host, Two ACSs, One 9310 LSM, One SL8500 Library Configuration.....	B-2
B-2	One Host, One ACS, One 9310 LSM, Dual LMU Configuration	B-12
B-3	Two Host, One ACS, Two 9310 LSM Configuration	B-14
D-1	Dual IP Connection with two SL8500 Connections	D-3
D-2	Dual IP Configuration Sample - Different Subnetworks	D-6
D-3	SLConsole Display for SL8500 Ports 2A and 2B.....	D-9
D-4	Dual IP (one host IP) with Two SL8500 Network Connections	D-10
D-5	Dual IP (one host IP) with Two SL8500 Network Connections	D-12
D-6	SLConsole Display for SL8500 Ports 2A and 2B.....	D-14
D-7	Four IP Addresses (four host IPs) with Four SL8500 Network Connections.....	D-16
D-8	Four IP (four host IPs) with One Single Redundant Pair SL8500 Network Connections	D-18
D-9	Four IP (four host IPs) with One Dual Redundant Pair SL8500 Network Connections	D-20
D-10	Four IP (four host IPs) with One Dual Redundant Pair SL8500 and 2 Dual TCP/IP Network Connections	D-22
D-11	Example SLC Console (Monitor Trace)	D-23
E-1	Dual IP with Two SL3000 Network Connections - Different Subnetworks.....	E-3
E-2	Dual IP Configuration Sample - Different Subnetworks	E-6

E-3	SLConsole Display of SL3000 Ports 2A/2B IP Addresses	E-9
E-4	Dual IP (one host IP) with Two SL3000 Network Connections	E-10
E-5	Dual IP (one host IP) with Two SL3000 Network Connections	E-12
E-6	SLConsole Display of SL3000 Ports 2A/2B IP Addresses	E-14
F-1	VSM4 with 32 Ports	F-2
F-2	VSM4 with 16 Host Ports, 16 RTD Ports	F-6
F-3	VSM4 with 20 Host Ports, 12 RTD Ports	F-7
F-4	VSM4 with 16 Ports	F-8
F-5	VSM4 with 8 Host Ports, 8 RTD Ports	F-10
F-6	VSM4 with 10 Host Ports, 6 RTD Ports	F-11
F-7	Configuration Diagram: Single MVS Host Connected to a VSM4 Using ESCON Directors... F-12	
F-8	Logical Paths for VSM 4 with 32 Ports, 31 Hosts, 16 RTDs	F-14
G-1	VSM4 with 6 ICE Cards, 2 VCF Cards	G-2
G-2	VSM4 with 4 ICE Cards, 4 VCF Cards	G-2
G-3	VSM4 with 2 ICE Cards, 6 VCF Cards	G-3
G-4	VSM4 with 8 VCF Cards	G-3
G-5	VSM4 with 8 VCF Cards, FICON Directors for 16 RTDs	G-7
G-6	VSM4 with 8 VCF Cards, 8 Host Ports, FICON Directors for 8 RTDs, 4 CLINK Ports ...	G-8
G-7	Dual ACS Bi-Directional Clustered VTSS Configuration	G-9
G-8	Configuration Diagram: Single MVS Host Connected to a VSM4 Using FICON Directors.... G-11	
H-1	VSM5 with 8 VCF Cards - Max 16 RTDs	H-2
H-2	VSM5 with 6 VCF Cards, 2 Empty Card Slots	H-2
H-3	VSM5 with 4 VCF Cards, 4 Empty Card Slots	H-3
H-4	VSM5 with 8 VCF cards - Max 32 RTDs	H-4
H-5	VSM5 with 8 VCF Cards, FICON Directors for 32 RTDs	H-8
H-6	Channel Configuration Status Screen	H-10
H-7	Channel Configuration and RTD Path Validation Screen	H-11
H-8	Real Tape Drive Status Screen.....	H-12
H-9	Configuration Diagram: Single MVS Host Connected to a VSM5 Using FICON Directors.... H-13	
I-1	Supported Card Configurations for VSM5 ESCON/FICON	I-2
I-2	VSM5 with 4 ICE cards, 4 VCF cards	I-3
I-3	VSM5 with 16 Host Ports, 16 RTD Ports	I-6
I-4	Configuration Diagram: Single MVS Host Connected to a VSM5 Using ESCON Directors... I-7	
K-1	Configuration Diagram: Single MVS Host Connected to a VSM 6 Using FICON Directors... K-6	
L-1	Tapeless VSM	L-2
L-2	Mixed Tapeless VSM	L-4
M-1	Relationship Between oVTCS Client and oVTCS CDS DB Server.....	M-1

List of Tables

1-1	HSC/VTCS Configuration Checklist	1-1
1-2	CDS Levels for Supported VTCS Versions	1-17
1-3	Valid CONFIG CDSLEVEL Values for VTCS 7.3.....	1-17
1-4	Log File Sizing	1-19
1-5	RTD Model/MVC Media Values	1-23
2-1	Enable CDS Logging.....	2-1
2-2	Locked Resource Policy	2-2
2-3	AMT Settings	2-3
2-4	VTSS Model Maximum VTVs and Automigration Thresholds	2-3
2-5	Delete Scratched VTVs Settings	2-4
2-6	VTV Page Size.....	2-5
2-7	Maximum VTV Size.....	2-6
2-8	Maximum RTDS per VTSS - 16 or 32	2-7
2-9	RTD/CLINK Addresses - Maximum 32 RTDs.....	2-8
2-10	Maximum VTVs per MVC.....	2-11
2-11	Hosts Disabled from Migration, Consolidation, or Export by VTV or MGMTclas	2-11
2-12	Recall VTVs with Read Data Checks.....	2-12
2-13	ETTFB Policy.....	2-12
2-14	Stacked Migrates Policy	2-13
2-15	Stacked Migrates Policy	2-15
2-16	Minimum VTV Residency Interval	2-16
2-17	Maximum Residency Interval before Deletion.....	2-17
2-18	Immediate Migration Delay Interval	2-17
2-19	Maximum and Minimum Concurrent Migration Tasks	2-18
2-20	Prioritizing Automatic and Immediate Migration Workloads	2-19
2-21	Controlling Migration Workloads.....	2-19
2-22	Conditionally Change Migration Preferencing/Depreferencing	2-20
2-23	Dependent Immediate Migration	2-21
2-24	VTV Replication - Always or Changed VTV Only	2-25
2-25	VTV Replication - Synchronous or Asynchronous	2-26
2-26	VTV Enhanced Replication.....	2-27
2-27	Standard Reclaim Threshold	2-30
2-28	Dynamic Reclaim Threshold	2-32
2-29	Dynamic Reclaim - Data Movement	2-34
2-30	VLE Reclaim Threshold	2-38
2-31	Free MVCs Threshold.....	2-38
2-32	Eligible/Total MVCs Threshold	2-39
2-33	Maximum MVCs Processed Per Reclaim	2-40
2-34	Maximum MVCs Concurrently Processed for Reclamation and Drain	2-41
2-35	Hosts Disabled from Reclamation.....	2-41
2-36	MVC Retain Interval.....	2-41
2-37	MVC Mount Timeout Interval.....	2-42
2-38	Non-Library Resident MVC Policies.....	2-42
2-39	MVC Initialization on First Mount.....	2-43
4-1	LIBGEN Macros	4-2
6-1	VTCS Configuration Tasks	6-1
6-2	Security Class, Resource Class, and Access Values for MVC Pool Volser Authority.....	6-5
7-1	HSC Initialization.....	7-2
A-1	SLIDRIVS Addresses	A-6
A-2	SL8500 Drive Panel Layout Example - Outside Library	A-7
A-3	SL8500 Drive Panel Layout Example - Outside Library	A-8
A-4	Mapping Command Prefix Codes to Characters	A-12
D-1	Network Entries Worksheet	D-3

D-2	Network Entries Worksheet	D-10
E-1	Network Entries Worksheet	E-3
E-2	Network Entries Worksheet	E-10
F-1	VSM3 to VSM4 Comparison: Software and System Configuration ESCON Enhancements . F-1	
F-2	Optimizing VSM4 Port Operations	F-3
F-3	VSM4 Configuration Options - 32 Ports	F-4
F-4	VSM4 Configuration Options - 16 Ports	F-9
G-1	Supported Card Configurations for VSM4 FICON Front-End plus Back-End Connectivity. G-1	
G-2	Optimizing VSM4 FICON/ESCON Port Operations	G-5
H-1	VSM5 Features.....	H-1
H-2	Optimizing VSM5 FICON Port Operations	H-5
H-3	32 RTDs/Stacked Migrates Support Requirements for VTCS/NCS 7.0.....	H-6
I-1	Supported Card Configurations for VSM4 FICON Front-End plus Back-End Connectivity. I-1	
I-2	Optimizing VSM5 FICON/ESCON Port Operations	I-5
J-1	VSM5 New Model Capacities	J-1
J-2	VSM5 New Model Supported Channel Card Configurations	J-2
K-1	ELS Features	K-1
K-2	ELS Supporting PTFs for VLE.....	K-3
K-3	Add Channel Path Values	K-9
K-4	Add Control Unit Values	K-10
K-5	Add Devices Values.....	K-12

Preface

Oracle's StorageTek Enterprise Library Software (ELS) is a solution consisting of the following base software:

- StorageTek Storage Management Component (SMC)
- StorageTek Host Software Component (HSC)
- StorageTek Virtual Tape Control Software (VTCS)
- StorageTek Concurrent Disaster Recovery Test (CDRT)

Additionally, the following software is provided with the ELS package:

- StorageTek Library Content Manager (LCM) (formerly ExLM). LCM includes an enhanced version of the product formerly known as Offsite Vault Feature.
- StorageTek Client System Component for MVS Environments (MVS/CSC)
- StorageTek LibraryStation

Audience

This guide is for StorageTek or customer personnel who are responsible for configuring HSC and VTCS. See *Introducing ELS* for an overview of the ELS solution. See *ELS Command, Control Statement, and Utility Reference* for information about VTCS, SMC, and HSC commands, control statements, and utilities. See *Configuring and Managing SMC* for information about configuring SMC. See *Installing ELS* for information about installing ELS.

Prerequisites

To perform the tasks described in this guide, you should understand the following:

- z/OS operating system
- JES2 or JES3
- Enterprise Library Software (ELS)

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at

<http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc>.

Access to Oracle Support

Oracle customers that have purchased support have access to electronic support through My Oracle Support. For information, visit

<http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> or visit

<http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> if you are hearing impaired.

What's New?

This revision includes the following updates:

- Added the TCP host-to-host communications method, designed to ease burden on HSC system resources.

See the following:

- ["HSC CDS Performance and Sharing Tips"](#)
 - ["Host-to-Host Communication Options"](#)
 - ["Adding Definitions for ACF/VTAM Communications"](#)
 - ["Creating the HSC PARMLIB Member"](#)
- Added support for IBM LTO Generation 7 and Generation 8 tape drives.
See ["Using the STORclas MEDIA Parameter for MVC Media Preferencing"](#) for more information.
 - General maintenance updates to address code fixes.

Planning for HSC/VTCS Configuration

Table 1–1 is designed to help plan and verify completion of your system’s configuration tasks, and if you look at the notes, you will see that, depending on your situation (new or upgrade install, adding hardware or not, and so forth), you may not have to do anything for a specific task except check it off.

Table 1–1 HSC/VTCS Configuration Checklist

Task	Required or Optional?	Notes	Check to Verify Completion
"Determining HSC/VTCS Configuration Values"	Required	Plan configuration values because you cannot take defaults here.	
"Planning VTCS Operating Policies"	Optional	You have a choice, that is, you can take the defaults and change them later. If you take that route, it is not a bad idea to take a high level pass through this chapter with and read it thoroughly when you have the time.	
"Defining A Security System User ID for HSC and VTCS" "Configuring MVS Device Numbers and Esoterics" "Setting the MVS Missing Interrupt Handler (MIH) Value"	Required	None of these tasks is difficult, but they are all critical. For example, depending on the default settings of your security system, VSM may not be able to mount and write to MVCs until you have defined a security system user ID for HSC and tape volume security profiles for the MVCs.	
"Creating the HSC CDS"	Required	Define the ACSs, LSMs, and drives so ELS can use them.	
"Defining Volumes to ELS"	Required	ELS 7.3 lets you define all volumes: Library Volumes, MVCs, VTVs, and cleaning cartridge, and you do this one time only with the POOLPARM/VOLPARM statements and SET VOLPARM utility.	
"Adding Definitions for ACF/VTAM Communications"	Required	If you set COMMPATH METHOD to VTAM, you must define the appropriate values to VTAM.	

Table 1–1 (Cont.) HSC/VTCS Configuration Checklist

Task	Required or Optional?	Notes	Check to Verify Completion
"Defining the SYS1.PARMLIB Member SMFPRMxx"	Optional	This task is optional but highly recommended , because you need the SMF information to know how your system is performing.	
"Creating the HSC PARMLIB Member"	Required	The HSC PARMLIB is where critical items such as the COMMPATH and FEATURES statements reside.	
"Defining Command Authority"	Required	This is required to ensure that the correct personnel and applications have access to the ELS resources required.	
"Updating HSM"	Required	This is required if you are an HSM user and are routing HSM jobs to VSM.	
"Creating and Cataloging the HSC Startup Procedure"	Required	An HSC startup proc is required.	
"Reconfiguring a TapePlex"	Optional	How much of reconfiguring a TapePlex you actually have to do depends on whether you are adding any hardware or changing values in the PARMLIB data sets.	
"Building a Simple CONFIG Deck"	Required if you license VTCS	A step procedure to create a basic CONFIG deck	
"Updating the Tape Management System for VSM"	Required if you license VTCS	This is required to give your TMS access to VTVs and keep the TMS from stepping on MVCs.	
"Defining MVC Pool Volser Authority"	Required if you license VTCS	Updating your security package to give the HSC user update authority to MVCs.	
"Storing VTCS Locks in a Coupling Facility (Optional)"	Optional if you license VTCS	Recommended in some instances for performance reasons.	
"Configuring VTCS for 512 VTD Support"	Required if you license VTCS	A step procedure to configure 512 VTDs.	
"Starting and Stopping HSC"	Required	NA	

How VSM Measures Sizes and Capacities

VTCS uses the binary standard rather than the decimal standard in displaying and calculating sizes and capacities for VTVs and MVCs. Thus:

- 1 kilobyte(KB)=1024 bytes
- 1 megabyte(MB)=1000 kilobytes or 1000*1024 byte
- 1 gigabyte(GB)=1000 megabytes or 1000*1000*1024 bytes

VTCS uses the decimal standard in displaying and calculating sizes and capacities for VTSSs. Thus:

- 1 kilobyte(KB)=1000 bytes
- 1 megabyte(MB)=1000 kilobytes or 1000*1000 bytes
- 1 gigabyte(GB)=1000 megabytes or 1000*1000*1000 bytes

Determining HSC/VTCS Configuration Values

The following sections tell how to determine configuration values for the HSC/VTCS system. Many of the following sections provide configuration planning information for HSC followed by planning information for VTCS. For example, to define RTDs, you must first define them as library-attached transports, then as RTDs. If you want transports to function as "native Nearline only," you define them as library-attached transports only, not as RTDs. The following sections are ordered as "both HSC and VTCS" followed by "VTCS only."

Note: If you are configuring devices for VTCS, unless otherwise noted, in each of the following sections, the values you determine must match wherever you use them. For example, the unit addresses described in "[Planning for Library-Attached Transports](#)" must match the following:

- The device addresses on the HSC `SLIDRIVS` macro.
 - The MVS device addresses you assign using the HCD facility (if you will share these transports with MVS).
-
-

Control Data Sets

Control data sets (CDSs) contain:

- inventory information on all volumes (real, virtual, and vaulted)
- configuration information, including how many ACSs, LSMs, tape transports, and so forth
- information about library hardware resource ownership across multiple processors

The control data sets are:

Primary CDS (required)

The primary CDS resides on a DASD. The primary CDS must be accessible by all CPUs configured with the library. All configuration information about the TapePlex is stored in this data set, and it is continuously updated to reflect changes in the library.

Secondary CDS (optional, but strongly recommended)

This data set is an exact duplicate of the primary CDS and is also continuously updated to reflect changes in the library. In general, if the primary CDS becomes corrupted, the HSC continues to operate by automatically switching to the secondary CDS. The secondary CDS then becomes the new primary CDS, but the data set name remains unchanged.

Standby CDS (optional, but strongly recommended)

This data set is a formatted CDS containing only the first CDS record and cannot be used as is to continue normal operations. In general, if the primary CDS becomes corrupted, and a switch to the secondary occurs, the contents of the secondary CDS are copied to the standby. The standby CDS then becomes the new secondary data set, but the data set name remains unchanged.

Control Data Set Placement

For performance and recovery considerations, each copy of the CDS should reside on a separate disk volume. Separate control units and channel paths are also recommended to provide adequate recovery conditions.

If possible, place each control data set on its own dedicated volume. Do not place HSC control data sets on any volumes containing HSC journal data sets.

The following restrictions also apply:

- VSM **does not** support copies of the CDS at multiple sites (for example, Primary CDS at one site and Secondary at another). A link failure would allow the two sites to run independently, and VSM cannot enforce separation of all resources. This prevents reconciliation of the two divergent CDSs as can be accomplished in a pure non-VSM environment.
- Similarly, copies of the **entire** CDS at two sites where a link failure may occur is not recommended. Refer to the *ELS Legacy Interfaces Reference* for more information.
- Copies of the **entire** CDS at two sites that are **not** linked is not allowed.
- The client-server approach of using HSC on only one host and SMC on all other hosts is recommended for TapePlexes spanning multiple geographic locations.

Control Data Set Sharing

The following sections details sharing requirements and recommendations.

Sharing Requirements

Sharing CDS devices between multiple host systems and processors requires that the CDS devices be defined with the `SHARED=YES` attribute using the IBM Hardware Configuration Definition (HCD) facility.

Note: Unit Control Blocks (UCBs) for HSC CDS volumes can be defined either as `LOCANY=NO` or `LOCANY=YES` using the IBM Hardware Configuration Definition (HCD) facility.

If the control data sets are required by more than one system, then the primary, secondary, and standby control data sets must be capable of being shared in read/write mode by all host systems that share the devices.

Sharing Recommendations

Several HSC functions use the MVS RESERVE/RELEASE service to serialize CDS access between multiple hosts. If your MVS system is configured to issue device reserves in support of the RESERVE/RELEASE service, I/O contention to the device can be noticed from other systems. This is because the device will appear busy to any system not holding the reserve, and cause those I/O operations to be queued. If queued long enough, IOS start pending messages will begin to appear on those systems.

The dynamic nature of the MVS operating environment makes it impossible to predict a maximum length of time that the HSC, or any other authorized program, will hold a RESERVE. The vast majority of HSC functions hold a RESERVE for a short interval, generally only a few I/O operations to the CDS. However, some functions require a longer interval to perform the work. For example, the BACKUP utility can take several

minutes to finish; the RESERVE must be held for the duration to guarantee the integrity of the backup.

Given the above, consider these recommendations:

- Do **not** place copies of the HSC control data set on the same volume(s) as other data sets that generate high I/O activity to that volume during normal processing. This applies to all control data set copies including the secondary (shadow) and the standby. If possible, place each control data set on its own dedicated volume (no other data sets on the volume)
- Do **not** place other datasets that require RESERVE/RELEASE functions on the same volume as the HSC control data sets. You must ensure programs that interact with each other do not access multiple CDS volumes concurrently, or a deadly embrace scenario can occur.

For example, TMS catalog and audit data sets, JES Checkpoint, RACF databases, and DFSMSHsm journals are known to cause contention and lockout problems when on the same volume as HSC CDSs. A backup copy of a data set that is used only occasionally for restore purposes normally does not cause a significant problem. However, if response problems or lockouts do occur, an examination should be made of all ENQ and I/O activity related to that volume.

- The default or recommended missing interrupt handler time interval for most disk devices should be high enough to avoid problems with the HSC reserves. However, if I/O contention due to HSC reserves begins to cause significant problems, consider increasing the missing interrupt handler time interval. Specific recommendations are detailed in the following sections.
- Do **not** use the I/O Timing facility (the MIH IOTIMING parameter in the IECIOSxx member of the system PARMLIB on MVS systems) for devices containing the HSC control data sets.
- Do **not** use the RECOVERY statement (in the IECIOSxx member) for devices containing the HSC control data sets.
- Do **not** specify IOTHSWAP=YES (in the IECIOSxx member) for devices containing the HSC control data sets.

Additional Recommendations When HSC Runs on an MVS Guest Under VM

In order to prevent errors caused by contention lockout situations with other hosts when an MVS system is running as a guest under VM, it is strongly recommended that the VM missing interrupt interval for each disk device containing the primary, secondary, and standby control data sets be set to a value **greater** (at least 1 minute) than the value of the MVS missing interrupt interval.

Note: If using SLUADMIN utility to backup and restore the CDS, set the MVS missing interrupt interval value higher than the time it usually takes to complete the backup or restore.

The VM missing interrupt interval is specified on the SET MITIME command or on the INIT_MITIME statement. The MVS missing interrupt interval is specified on the TIME parameter of the MIH statement. Refer to the appropriate IBM publications for more information about setting these values.

Additional Recommendations When the CDS is Shared Between MVS and VM

When there are no MVS guest systems running under VM, but only "native" MVS systems that share the CDS with VM systems, the MVS and VM missing interrupt intervals act independently of each other, however, for long-term ease of management, you may want to follow the same recommendations for when HSC runs on an MVS guest under VM in case you ever change a native MVS system to run as an MVS guest under VM:

- It is **strongly** recommended that the VM missing interrupt interval for each disk device containing the primary, secondary, and standby control data sets be set to a value greater (at least 1 minute) than the value of the MVS missing interrupt interval.
- If you are using the SLUADMIN utility to backup and restore the CDS, set the MVS missing interrupt interval value higher than the time it usually takes to complete the backup or restore.

The VM missing interrupt interval is specified on the SET MITIME command or on the INIT_MITIME statement. The MVS missing interrupt interval is specified on the TIME parameter of the MIH statement. Refer to the appropriate IBM publications for more information about setting these values.

Other considerations that apply to sharing a CDS between MVS and VM systems include, but are not limited to, the following:

- DFSMSdss DEFRAG running on MVS can cause lockout conditions on VM. Also, because DEFRAG can move data sets, control data set integrity can be compromised. Avoid running a DEFRAG on CDS volumes while the HSC is in operation on any shared system.
- DFSMSHsm processing of CDS volumes has also been known to cause lockout conditions. If possible, avoid DFSMSHsm COMPRESS, COPY, DEFRAG, DUMP, PRINT, RELEASE, and RESTORE operations on CDS volumes during peak periods of HSC operation.

Make sure the HSC command prefix does not conflict with any of the VM facilities, such as the CP line editing symbol.

Limiting CDS Access Privileges

Use your system security software (such as RACF, ACF-2, or TopSecret) to give WRITE access to the CDS primary, secondary and standby data sets for the user ID associated with the HSC started task. You must also submit SLUADMIN utility jobs with this user ID because some SLUADMIN utility functions update the CDS.

Caution: DO NOT ASSIGN CDS WRITE ACCESS TO ANY OTHER USER.

ELS utilities such as SLUCONDB and ExPR require READ access to the CDS data sets, as LCM does. Assign CDS READ access only to user IDs that require this access.

Serializing CDSs

Resource serialization in the IBM z/OS operating system is generally accomplished using either the IBM Global Resource Serialization (GRS) facility or a third-party software product such as Unicenter CA-MIM/MII.

A resource is identified by two names:

- An 8-character QNAME
- A 1-255 character RNAME

The relationship between the two names is hierarchical; different RNAMEs can be used to describe specific different resources that are generically related to a single QNAME.

The HSC and HSC utilities use two principal QNAMES:

- The default QNAME of STKALSQL can be changed and is used by the HSC to serialize access to the HSC CDS and by LCM to serialize LSM-related processing.

Note: If you have changed STKALSQL to a different QNAME, just substitute your name for the STKALSQL references in this documentation.

- The StorageTek-defined QNAME of STKENQNM cannot be changed and is used to serialize HSC utilities, such as the SET utility.

Set Utility Serialization

One of the facilities provided by the SET utility is to change the QNAME that was defined during LIBGEN (the STKALSQL default or customer-defined QNAME) and stored in the CDS. The SET utility and HSC also use the STKENQNM QNAME to maintain serialization while the customer-defined QNAME is being changed.

The SET utility issues two RESERVEs against the CDS prior to an update:

- a RESERVE with the StorageTek-defined QNAME "STKENQNM"
- a RESERVE using the existing customer-defined QNAME (or the default value of "STKALSQL")

When the HSC is started on any host, it initially serializes on the CDS using the StorageTek-defined QNAME. This prevents the HSC from being started while the customer-defined QNAME is potentially in the process of being changed. If the serialization is successful (no SET utility in progress), the customer-defined QNAME is read from the CDS and is used for future serialization requests.

GRS Serialization

In GRS environments only, resources can be serialized across multiple systems by global (SCOPE=SYSTEMS) requests or serialized on individual systems only by local (SCOPE=SYSTEM) requests. GRS provides three ways (Resource Name Lists) to change how a request is to be processed:

- The Systems Exclusion Resource Name List (referred to as the Systems Exclusion List) allows global requests to be converted to local requests.
- The Systems Inclusion Resource Name List allows local requests to be converted to global requests.
- The Resource Conversion Resource Name List (referred to as the Reserve Conversion List) allows RESERVE requests to be suppressed.

These three lists are built from RNLDEF statements specified in the GRSRNLxx member of the system PARMLIB in GRS environments.

Note when a RESERVE is issued, there is also an enqueue associated with the RESERVE. If a matching entry is found in the Systems Exclusion List, the RESERVE is

issued, and the enqueue associated with a RESERVE is issued as a local enqueue request. If a matching entry is not found in the Systems Exclusion List, the RESERVE is issued and the enqueue associated with a RESERVE is issued as global enqueue request.

Caution: If no matching entry is found in the Systems Exclusion List and the Reserve Conversion List, double serialization occurs. Avoid this at all costs. The *IBM z/OS MVS Planning: Global Resource Serialization* publication shows a diagram of this process.

Multiple HSC TapePlex Considerations

If you have multiple HSC TapePlexes (each TapePlex must use a different set of primary, secondary, and standby CDSs) within the same GRS or MIM/MII serialization complex, make sure you change the default HSC QNAME STKALSQN to a different value for each HSC TapePlex. This ensures the serialization resources for one TapePlex does not delay the serialization of resources for the other TapePlexes. The default name can be changed either through the MAJNAME parameter on the SLILIBRY macro in LIBGEN or through the SET MAJNAME command of the SLUADMIN utility.

Note: With multiple TapePlexes, remember to replicate the STKALSQN examples shown in this documentation as many times as necessary and change STKALSQN to the different values you chose for each TapePlex.

Example

For two HSC TapePlexes, you could change the default HSC QNAME of STKALSQN to HSCPLEX1 for one TapePlex and HSCPLEX2 for the other. This allows the two TapePlexes to operate simultaneously without interfering with each other.

As a specific case, LCM management runs issue global enqueues for the HSC QNAME and an RNAME of PROCESSLSMaa:ll, where aa is the ACSid and ll is the LSMid. In a configuration of two HSC TapePlexes, there is an LSM 00 in an ACS 00 for each TapePlex, which results in the same RNAME of PROCESSLSM00:00. Two simultaneous LCM management runs for LSM 00 in ACS 00. Two simultaneous LCM management runs for LSM 00 in ACS 00 will conflict unless the HSC QNAMEs are different for the two TapePlexes.

HSC CDS Performance and Sharing Tips

Follow these tips for optimal CDS performance and sharing:

- Place each copy of the HSC CDS (primary, secondary, and standby) on its own dedicated volume with no other data sets on those volumes. This is especially important when you do not convert HSC RESERVEs. A RESERVE issued by one system locks out access to all data on the volume for all other systems, so if you put a catalog or the TMC on the same volume as an HSC CDS, you will create a performance problem. Isolating the HSC CDSs to dedicated volumes can also simplify recovery after an outage.
- Make sure all HSC CDS devices are genned as SHARED in your I/O device definitions. To do this, assign a value of YES to the SHARED feature for the device on the HCD View Device Parameter/Feature Definition panel and then activate the Input/Output Definition File (IODF) containing the change.

- Make sure you use the TCP communications method for all HSC host-to-host communication. This is specified by the METHOD parameter on the COMMPATH command and control statement.
- If you use both HSC and VTCS software, and all hosts are running on release 6.1 or later, consider converting the CDS to the “F” level format (or later), which reduces CDS I/O during HSC/VTCS initialization and periodically thereafter when VTCS refreshes its cache.
- If you use VTCS in a sysplex configuration and you believe you have a CDS performance problem because of VTCS, contact StorageTek Software Support to have the problem analyzed before implementing VTCS lock data in a Coupling Facility structure.

Unicenter CA-MIM/MII Considerations

Start the HSC only after the CA-MII address space has fully initialized.

Follow Computer Associates’ recommendations to start the CA-MII address space using their MIMASC utility, make sure the CA-MII address space is in the WLM SYSTEM service class so that it runs with a dispatching priority of FF, and add a PPT entry in the SYS1.PARMLIB(SCHEDxx) member on all systems in the MIM/MII serialization complex for the MIMDRBGN program. Refer to the CA documentation for other tuning recommendations.

Whenever you need to add a QNAME statement for the HSC, you must specify SCOPE=SYSTEMS and not SCOPE=RESERVES on the QNAME statement, since other NCS products (LCM and LibraryStation) issue global enqueues that must be propagated to all systems. Specifying SCOPE=RESERVES prevents these enqueues from being propagated and will cause problems.

GRS Considerations

For GRS ring configurations, the ACCELSYS and RESMIL parameters (specified on GRSDEF statements in the GRSCNFxx member of the system PARMLIB) can influence ring performance. Refer to *IBM z/OS MVS Planning: Global Resource Serialization* for information about how best to set these values.

When HSC RESERVEs Must Remain as RESERVEs (All Environments)

Caution: Do not convert HSC RESERVEs to global enqueues in the following configurations and environments, or CDS data integrity will be compromised and CDS damage will occur.

- No serialization facility or product (e.g., GRS or CA-MIM/MII) is used, and the HSC CDSs are on devices that are shared by multiple systems.
 - Even unconverted RESERVEs are insufficient to guarantee data integrity in this case. You should consider why you are trying to share data without a serialization product.
- A serialization facility or product is used, but the HSC CDSs are on devices shared by multiple serialization complexes (GRS or MIM/MII).
 - Even unconverted RESERVEs are insufficient to guarantee data integrity in this case. The *IBM z/OS MVS Planning: Global Resource Serialization* publication indicates that multiple serialization complexes cannot share resources.

- Computer Associates cites this restriction in their paper *Unicenter CA-MII Data Sharing for z/OS*.
- The environment consists of a GRS star configuration that attempts to share the HSC CDSs between a sysplex and a system that is not part of the sysplex.
 - IBM cites this restriction in *z/OS MVS Planning: Global Resource Serialization*.
- The HSC CDSs reside on devices shared by z/OS and z/VM systems.
 - IBM cites this restriction in *z/OS MVS Planning: Global Resource Serialization*.
 - Computer Associates cites this restriction in their paper *Unicenter CA-MII Data Sharing for z/OS*.

When You May Want to Leave HSC RESERVEs as RESERVEs

You may not want to convert HSC RESERVEs to global enqueues in the following configurations and environments.

- You use GRS in a ring configuration and converting HSC RESERVEs affects ring performance.
 - IBM cites this consideration in *z/OS MVS Planning: Global Resource Serialization*.
 - Performance can suffer due to propagation delays in GRS ring configurations. Deciding whether or not to convert RESERVEs is best determined by actual experience, but as an arbitrary rule, do not convert RESERVEs if there are more than three or four systems in your GRS ring.
- You use GRS in a ring configuration and converting HSC RESERVEs affects HSC/VTCS performance.
 - Because of ring propagation delays for global enqueues, HSC and VTCS throughput may significantly degrade during periods of high HSC and/or VTCS activity.
- You have a large virtual tape configuration. For example, if you have several million Virtual Tape Volumes (VTVs) defined and have not yet migrated the CDS to the “F” level format, you may experience slow HSC/VTCS initialization times and higher CDS I/O activity due to the need for VTCS to initialize and periodically refresh its cache of VTV information.
- You use GDPS HyperSwap, which requires specifying a pattern RNL to convert all RESERVEs, but you do not want to convert the HSC RESERVEs. These competing goals can be accommodated in a GRS environment by performing the following tasks:
 - Ensure that no device defined to GDPS can have a RESERVE issued to it.
 - Place the HSC CDSs on devices that are outside the scope of GDPS and HyperSwap control (i.e., not defined to GDPS) and therefore are not HyperSwap eligible.
 - Specify the following RNLDEF statements in the GRNRNLxx member of the system.PARMLIB on all systems:

```
RNLDEF RNL(EXCL) TYPE(GENERIC) QNAME(STKALSQN)
RNAME(hsc.primarycds.datasetname)
RNLDEF RNL(EXCL) TYPE(GENERIC) QNAME(STKALSQN)
RNAME(hsc.secondarycds.datasetname)
RNLDEF RNL(EXCL) TYPE(GENERIC) QNAME(STKALSQN)
RNAME(hsc.standbycds.datasetname)
RNLDEF RNL(CON) TYPE(GENERIC) QNAME(*)
```

The Reserve Conversion List is not searched if an entry is found in the Systems Exclusion List, thus GDPS is unaware of the RESERVEs processed by GRS. The net effect is that the GDPS HyperSwap requirement of a pattern RNL to convert all RESERVEs is satisfied, but the HSC RESERVEs are not converted.

When HSC RESERVEs Must be Converted to Global Enqueues

Convert HSC RESERVEs to global enqueues if the HSC CDSs are on devices that are defined to GDPS.

How to Leave HSC RESERVEs as RESERVEs in a GRS Environment

If you do not convert HSC RESERVEs, remove all RNLDEF RNL(CON) statements for the HSC from the GRSRNLxx member on all systems. Even though the Reserve Conversion List is not searched if an entry is found for the HSC RESERVE in the Systems Exclusion List, it is good practice to keep your GRSRNLxx definitions up to date.

More importantly, if there are no RNL(EXCL) or RNL(CON) statements at all for the HSC, both a SYSTEMS ENQ and a RESERVE are issued for the HSC CDS. This results in double serialization and reduced performance. Refer to *IBM z/OS MVS Planning: Global Resource Serialization* for a diagram showing how RESERVEs are processed by GRS.

To avoid double serialization, perform **only one** of the following tasks:

- Do not convert HSC RESERVEs by adding the recommended RNL(EXCL) statements for the HSC.
- Convert HSC RESERVEs by adding the recommended RNLDEF(CON) statements for the HSC.

To not convert HSC RESERVEs, include one RNL(EXCL) statement for the STKALSQN resource for each copy of the CDS (primary, secondary, and standby) and one statement for the STKENQNM resource as follows:

```
RNLDEF RNL(EXCL) TYPE(GENERIC) QNAME(STKALSQN) RNAME(hsc.primarycds.datasetname)
RNLDEF RNL(EXCL) TYPE(GENERIC) QNAME(STKALSQN)
RNAME(hsc.secondarycds.datasetname)
RNLDEF RNL(EXCL) TYPE(GENERIC) QNAME(STKALSQN) RNAME(hsc.standbycds.datasetname)
RNLDEF RNL(EXCL) TYPE(GENERIC) QNAME(STKENQNM)
```

- Do not specify TYPE(SPECIFIC) on the RNLDEF statements for the STKALSQN resource.

The HSC code uses an RNAME 44 bytes long with blank fill on the right, so a TYPE(SPECIFIC) statement as shown will not be matched and will not produce the desired results.

- Do not remove the RNAME parameter from the RNLDEF statements for the STKALSQN resource.

The RNLDEF RNL(EXCL) statements tell GRS a RESERVE for the HSC CDS should be issued as a RESERVE and a global (SCOPE=SYSTEMS) enqueue issued for STKALSQN (including the enqueue associated with the RESERVE) should be changed to a local (SCOPE=SYSTEM) enqueue. The RNAME parameter limits these changes in scope to the CDS only.

Using a generic QNAME statement without an RNAME parameter tells GRS to change **all** global enqueues for STKALSQN to local enqueues. Since other NCS

products (LCM and LibraryStation) issue global enqueues with the expectation that they will be propagated to other systems and not reduced in scope, using a generic QNAME statement for STKALSQN without an RNAME parameter will cause problems. The RNAME parameter must be explicitly specified for each copy of the CDS as in the example above. The STKENQNM resource can continue to be defined using a generic QNAME only.

How to Leave HSC RESERVEs as RESERVEs in an MIM/MII Environment

For PROCESS=SELECT and PROCESS=ALLSYSTEMS environments:

- Whenever you need to add a QNAME statement for the HSC, specify SCOPE=SYSTEMS and not SCOPE=RESERVEs on the QNAME statement, since other NCS products (LCM and LibraryStation) issue global enqueues that must be propagated to all systems. Specifying SCOPE=RESERVEs prevents enqueues from being propagated and causes problems.

For PROCESS=SELECT and PROCESS=ALLSYSTEMS environments with LCM and/or LibraryStation:

- To not convert the CDS RESERVEs but propagate the other STKALSQN global enqueue requests, use an Exempt List and specify the QNAME statements as follows:

```
STKALSQN EXEMPT=YES, GDIF=YES, RESERVEs=KEEP, SCOPE=SYSTEMS
STKENQNM EXEMPT=NO, GDIF=NO, RESERVEs=KEEP, SCOPE=SYSTEMS
```

- Make sure an Exempt List can be used. If the GDIINIT statement specifies EXEMPT=NONE, either change it to specify EXEMPT=membername, or remove the EXEMPT parameter to cause the default member name of GDIEXEMPT to be used.
- In the GDIEXEMPT member (or whatever you named it) of the MIM parameter data set, specify:

```
LOCAL QNAME=STKALSQN RNAME=hsc.primarycds.datasetname
LOCAL QNAME=STKALSQN RNAME=hsc.secondarycds.datasetname
LOCAL QNAME=STKALSQN RNAME=hsc.standbycds.datasetname
GLOBAL QNAME=STKALSQN
```

- Enter the following in the MIM commands member for every system in the MIM/MII serialization complex.

```
SET GDIF EXEMPTRESERVEs=YES
```

Caution: The same value of EXEMPTRESERVEs=YES must be specified and in effect on all systems in the MIM/MII serialization complex at the same time, otherwise, a data integrity exposure will exist for all shared data. If you have any questions about changing to EXEMPTRESERVEs=YES, contact Computer Associates before making the change.

For PROCESS=SELECT environments without LCM and LibraryStation:

- You have a choice. PROCESS=SELECT implies that only those resources that are explicitly defined with QNAME statements will be processed by CA-MIM/MII, so you can do one of the following:

- Make sure there are **no** QNAME statements or Exempt List entries for the HSC, and CA-MIM/MII will leave all HSC RESERVEs alone.
- Explicitly add and enable new QNAME statements for the HSC, and make sure that Exempt List entries do not exist for the HSC. If you choose this option, you must specify GDIF=NO, RESERVES=KEEP, and SCOPE=SYSTEMS on the QNAME statements for the HSC. As a precaution, also specify EXEMPT=NO on the QNAME statements to prevent the Exempt List from overriding any QNAME statement values. For example:

```
STKALSQN EXEMPT=NO, GDIF=NO, RESERVES=KEEP, SCOPE=SYSTEMS
STKENQNM EXEMPT=NO, GDIF=NO, RESERVES=KEEP, SCOPE=SYSTEMS
```

For PROCESS=ALLSYSTEMS environments without LCM and LibraryStation:

- Add and enable new QNAME statements for the HSC and make sure Exempt List entries do not exist for the HSC.

Specify GDIF=NO, RESERVES=KEEP and SCOPE=SYSTEMS on the QNAME statements for the HSC. As a precaution, also specify EXEMPT=NO on the QNAME statements to prevent the Exempt List from overriding any QNAME statement values. You must code your HSC QNAME statements as shown below, since the dynamic addition of HSC resources by PROCESS=ALLSYSTEMS inappropriately defaults to assign the EXEMPT=YES and GDIF=YES attributes. For example:

```
STKALSQN EXEMPT=NO, GDIF=NO, RESERVES=KEEP, SCOPE=SYSTEMS
STKENQNM EXEMPT=NO, GDIF=NO, RESERVES=KEEP, SCOPE=SYSTEMS
```

How to Convert HSC RESERVEs to Global Enqueues in a GRS Environment

For both GRS ring and star configurations:

- Include one RNL(CON) statement for the STKALSQN resource and one RNL(CON) statement for the STKENQNM resource:

```
RNLDEF RNL (CON) TYPE (GENERIC) QNAME (STKALSQN)
RNLDEF RNL (CON) TYPE (GENERIC) QNAME (STKENQNM)
```

- Remove all RNLDEF RNL(EXCL) statements for the HSC from the GRSRNLxx member on all systems, otherwise, the HSC RESERVEs will not be converted to global enqueues. The Reserve Conversion List is not searched if an entry is found for the HSC RESERVE in the Systems Exclusion List.

How to Convert HSC RESERVEs to Global Enqueues in an MIM/MII Environment

For both PROCESS=SELECT and PROCESS=ALLSYSTEMS environments:

- Specify GRSRNL=EXCLUDE in the IEASYSxx member of the system PARMLIB on all systems to prevent RNL processing by GRS.
- If any systems are in a sysplex, refer to "Sysplex Considerations" in the "Advanced Topics" chapter of the *MII System Programmer's Guide* for other IEASYSxx requirements.
- Whenever you need to add a QNAME statement for the HSC, specify SCOPE=SYSTEMS and not SCOPE=RESERVES on the QNAME statement since other NCS products (LCM and LibraryStation) issue global enqueues that must be propagated to all systems. Specifying SCOPE=RESERVES prevents these enqueues from being propagated and causes problems.

For PROCESS=SELECT environments:

- Specify GDIF=YES, RESERVES=CONVERT, and SCOPE=SYSTEMS on the QNAME statements for the HSC. Make sure Exempt List entries do not exist for the HSC, and as a precaution, also specify EXEMPT=NO to prevent the Exempt List from overriding any QNAME statement values. For example:

```
STKALSQN EXEMPT=NO, GDIF=YES, RESERVES=CONVERT, SCOPE=SYSTEMS
STKENQNM EXEMPT=NO, GDIF=YES, RESERVES=CONVERT, SCOPE=SYSTEMS
```

For PROCESS=ALLSYSTEMS environments:

- If the GDIINIT statement specifies RESERVES=KEEP, you must specify the same QNAME statements as required for PROCESS=SELECT environments since the dynamic addition of HSC resources by PROCESS=ALLSYSTEMS will not (by default) assign the RESERVES=CONVERT attribute needed to convert the RESERVE:

```
STKALSQN EXEMPT=NO, GDIF=YES, RESERVES=CONVERT, SCOPE=SYSTEMS
STKENQNM EXEMPT=NO, GDIF=YES, RESERVES=CONVERT, SCOPE=SYSTEMS
```

If the GDIINIT statement specifies or is defaulted to RESERVES=CONVERT, you do not need to specify any QNAME statements unless you have special requirements that conflict with the defaults assigned to dynamically added QNAMEs. See "Selecting the GDIF Processing Mode" in the *Unicenter CA-MII Systems Programming Guide* for the defaults assigned to dynamically added QNAMEs.

CDS DASD Space Requirements

This section tells how to calculate the CDS DASD space required for ELS. The space in the CDS supports the following areas:

- HSC - Library hardware configuration and library volumes
- VTCS - VSM hardware configuration and VSM volumes (VTV/MVC)
- HSC SET VOLPARM - VOLPARM input card images
- HSC SET VAULTVOL - Vault volumes
- HSC SET CDKLOCK - Locking service for Open Systems ACSAPI

SLICREAT creates the HSC area from the HSC LIBGEN process. The other areas are created dynamically. The space for the dynamic areas must be initially allocated by SLICREAT; however, if insufficient space is allocated by SLICREAT, you can use CDS EXPAND procedure to dynamically expand the CDS. Refer to *Managing HSC and VTCS* for more information.

If you are specifying multiple CDSs (SLSCNTL2, SLSSTBY), StorageTek recommends that you allocate the same amount of space (in blocks) for all your data sets when you define them.

Note: If the data sets are defined with different space allocations, the HSC uses the size of the smallest data set to determine the number of 4K blocks that it will use for the CDS. The additional space in the other CDS data sets, if any, will not be used by the HSC.

The difference in the space between the minimum space required (returned by SLICREAT) and the size of the smallest CDS data set is formatted as CDS free blocks. These free blocks are available to create the dynamic areas (VTCS, and so forth).

Calculating DASD Space Requirements for POOLPARAM/VOLPARAM

The HSC SET VOLPARAM utility uses CDS space to store the input card images. The number of blocks required for this function can be calculated by the following formula:

$$(\text{input} / 50)$$

where *input* is the number of input records in the SET VOLPARAM utility.

When the SET VOLPARAM utility is used to specify VTV and MVC volumes, the VTCS areas for these are built when the SET VOLPARAM data is applied. See the following section for information on calculating space for VTVs and MVCs.

Calculating DASD Space Requirements for VTCS

This section tells how to calculate the additional CDS DASD space required for VTCS. The additional number of 4K blocks required in the CDS for VTCS is expressed as:

- For F and G format CDSs:

$$(\# \text{ VTV ranges}) + (\# \text{ VTV ranges})/862 + (\# \text{ VTVs defined})/23 + (\# \text{ VTVs defined})/19826 + (\# \text{ MVC ranges}) + (\# \text{ MVCs defined})/37 + 18 * (\# \text{ of VTSSs}) + 14$$

- For H format CDSs:

$$(\# \text{ VTV ranges}) + (\# \text{ VTV ranges})/862 + (\# \text{ VTVs defined})/23 + (\# \text{ VTVs defined})/19826 + (\# \text{ MVC ranges}) + (\# \text{ MVCs defined})/37 + 18 * (\# \text{ of VTSSs}) + 14 + (\# \text{ MVC volumes to be used for Dynamic Reclaim})/7$$

- For I format CDSs:

- If dynamic reclaim is not being used, use the H format calculation.
- If dynamic reclaim is being used, use the H format calculation plus the following:

$$(\# \text{ of expected 32GB VTVs}/182) + (\# \text{ of expected 32GB VTVs} / 182/ 32024) + 1$$

then round up to the closest whole number.

Your Oracle representative will run the VSM pre-sales planning tool to identify VSM candidate data sets. This will assist you in estimating the number of VTVs and MVCs required.

Calculating DASD Space Requirements for VAULTVOL

The HSC VAULT facility uses the CDS to store information on Vaulted Volumes. The number of blocks required for this function can be calculated by the following formula:

$$(\text{nnnn} * 1.2 / 99)$$

where *nnnn* is the number of Vaulted Volumes specified in the SET VAULTVOL utility.

To increase/decrease the number of Vaulted Volumes, the CDS MERGE procedure must be run.

Calculating DASD Space Requirements for CDKLOCK

The HSC Open Systems locking service uses a CDS subfile to store resource lock information. This subfile is only required if the open system platforms use the XAPI enhancement to the ACSAPI to provide resource (drive or volume) serialization. The number of blocks required for this subfile can be calculated by the following formula:

$$(xxx/10) + (xxx/15) + (xxx/20)$$

where *xxx* is the maximum number of HSC drives that are used by the Open System platforms.

For example, if 200 HSC drives can be used by Open Systems, allocate 44 blocks for the CDKLOCK subfile. To increase or decrease the size of the CDKLOCK subfile follow the MERGEcds procedure.

CDS VTCS Level Requirements

Note that there are different CDS levels for HSC and VTCS, and some key VTCS functionality is only enabled by migrating to the required CDS level as described in [Table 1-2](#). You can determine what level CDS you currently have with the HSC D CDS command as shown in the following example. In this example:

- The HSC CDS level is 6.1

Note: The HSC CDS level is the same for HSC 6.1 and above, so D CDS shows CDS LEVEL = 060100 for all three releases.

- The HSC CDS level is the same for HSC 6.1 and above, so D CDS shows CDS LEVEL = 060100 for all three releases.

```
.SLS2716I Database Information 029 029
SYS00001 = VTCS.HARDWARE.CFG16Y.CDS
PRIVOL = ENG001      FLAGS(40) ACTIVE

JOURNALING NOT ACTIVE FOR THIS SUB-SYSTEM

LOG DATA SET STATUS:
SYS00013 = VTCS.HARDWARE.CFG16Y.HSCLOG1
ACTIVE
SYS00014 = VTCS.HARDWARE.CFG16Y.HSCLOG2
ACTIVE
UTILIZATION = .76%

CDS LEVEL = 060100      DATE = 20150313
CREATE     = I813156    TIME = 14:53:37
VSM CDS LEVEL = I

LAST CDS BACKUP = NONE
LAST CDS RESTORE = NONE
LAST NCO ON

ENQNAME = STKALSQN      - SMFTYPE = 255
CLEAN PREFIX = CLN      - LABTYPE = (00) SL
```



```
RECOVERY = (00) NO RECOVERY
THIS HOST IS = ECCL          - CDS BLOCK COUNT = 12,060
```

As described in [Table 1–2](#), each supported VTCS version supports only a subset of these VTCS Levels. If you are, therefore, running with multiple versions of VTCS against a CDS, it is important to ensure that the CDS is set at a level that is supported by all the versions being run. Also note that VTCS functions are available only by running the CDS at a certain level.

Table 1–2 CDS Levels for Supported VTCS Versions

VTCS CDS Level	Valid VTCS/NCS Versions	Enhancements
E	6.0, 6.1, 6.2, 7.0	<ul style="list-style-type: none"> ■ 4 MVC copies ■ 800 Mb VTVs
F	6.1, 6.2, 7.0, 7.1, 7.2, 7.3	<ul style="list-style-type: none"> ■ Near Continuous Operations (NCO) ■ Bi-directional clustering ■ Improved CDS I/O performances - reduces the I/O required to manage virtual scratch sub-pools
G	6.2, 7.0, 7.1, 7.2, 7.3	<ul style="list-style-type: none"> ■ 400Mb/800Mb/2Gb/4Gb VTVs ■ Standard/Large VTV Pages ■ 65000 VTVs per MVC
H	7.1, 7.2, 7.3	<ul style="list-style-type: none"> ■ Dynamic Reclaim ■ Autonomous Device Support
I	7.3	VSM 6 32 GB VTV support

Specifying Valid CDS Levels for VTCS 7.3

[Table 1–3](#) describes the valid CDS levels for VTCS 7.3 and the corresponding CONFIG CDSLEVEL values.

Table 1–3 Valid CONFIG CDSLEVEL Values for VTCS 7.3

CDS VTCS Level	CDSLEVEL Value
F	V61ABOVE
G	V62ABOVE
H	V71ABOVE
I	V73ABOVE

Note:

- VTCS 7.3 will tolerate sharing the CDS with down-level VTCS versions. However, to run mixed 7.3 and down-level you need to run at a CDS level valid for **all** versions. For example, VTCS 7.3 and VTCS 7.2 can share the same CDS, but the VTCS level **must** be H.
 - Conversely, if you want specific features, you **must** be at the required CDS level, which may require you to upgrade down-level hosts. For example, VSM 6 32 GB VTVs requires CDS Level I, which is only valid at VTCS 7.3 and above.
-

Backing Out From an I Level CDS

To downgrade from an I level CDS:

1. Run the SLUADMIN BACKUP command against the current CDS.
2. Enter the VTVRPT command and identify all 32GB VTVs.
3. Copy the 32GB VTVs to non-32GB VTVs.
4. Using the VTVMAINT command, change the management class for all 32GB VTVs to a management class that contains only DELSCR(YES).
5. Run the SLUADMIN SCRATCH utility to scratch the 32GB VTVs. At this point the 32GB VTVs are no longer recoverable.
6. Stop all active HSCs. See ["Stopping HSC."](#)
7. Update the VTCS configuration.
 - a. Enter CONFIG CDSLEVEL.
 - b. Enter CONFIG GLOBAL MAXVTVSZ.
 - c. Code CONFIG VTVVOL and CONFIG MVCVOL statements. Run a VTCS DECOMP utility to show the current VTVVOL and MVCVOL statements.
8. Run the SLUADMIN CONFIG utility.
9. Update any MGMTCLAS that references MAXVTVSZ(32000).
10. Start HSC. See ["Starting HSC."](#)

Using CDS Logging to Help Recover the CDS

The HSC Transaction Logging Service is an internal HSC service that records information into one or, optionally, two Log Files. If all copies of the CDS fail (that is, you lose the Primary, Secondary, and Standby copies), recovery from a backup can be time consuming. When recovering from a backup, CDS logging can aid in resynchronizing the CDS with the data content of the VSM subsystem.

Planning for CDS logging includes the following:

- Do you want one or two copies of the log file? Two copies gives more redundancy but takes up additional DASD space.
- How much DASD space do you need for logging? It is dependent on the activity on your system and the frequency with which you do backups. Log file space requirements depend on the number of VTVs mounted between log file off loads. Use [Table 1-4](#) to size your log file space allocation.

Table 1–4 Log File Sizing

Log file size (Mb)	VTV Mounts
1	10,000
2	20,000
4	40,000
8	80,000
16	160,000
32	320,000
64	640,000

- Where do you place the log files? Basically, each log file should reside on different volumes and not on the same volume as any of the CDS copies.

Planning for LSMs

Planning for LSMs is an important phase of the installation process. Effective planning for StorageTek automated tape libraries is part of a smooth ACS installation process. Ensure that planning considerations are reviewed for:

- Placement of ACS network connections on your network
- LSM Unit address requirements
- LSM/Pass-thru Port (PTP) relationships
- Host-to-Host communication options
- SL8500 libraries
- SL3000 libraries

Placement of ACS Network Connections on Your Network

TCP/IP Network attached libraries should be attached to a separate subnet or a controlled network to protect it from floods of ARP broadcasts. All Streamline libraries are TCP/IP network attached. In addition, some legacy LSM hardware can also be attached through the TCP/IP network.

LSM Unit Address Requirements

Some legacy hardware LSMs may use a LMU Station that emulates a 3278-2 terminal attached to a local controller. This LMU station must be assigned an MVS unit address. The LMU must be defined as a 3x74 local controller. The HCD facility is used to assign MVS unit addresses as a 3278-2 terminal to the LMU station.

LSM Pass-Thru Port (PTP) Relationships

If your ACS contains two or more connected LSMs, you must define the LSM/PTP relationships. The Master LSM controls the PTP; the Slave LSM does not. You define the Master/Slave LSM relationships in the SLILSM Macro. Legacy LSMs (e.g. 9310 Powderhorn) can be a master of a maximum of two PTPs. Even though a Legacy LSM may have more than two PTPs, it can control only two PTPs; it is possible for one LSM to be both a master and a slave. Streamline libraries are different. The SL8500 can control 8 PTPs. The SL3000 has no PTPs.

Host-to-Host Communication Options

To optimize performance, Oracle recommends that you set the HSC `COMMPATH METHOD` parameter to `TCP` to allow even sharing of resources in a multi-host configuration as shown in the example in ["Creating the HSC PARMLIB Member."](#)

Calculating Cartridge Capacity - SL8500 and SL3000

Message SLS0557I displays the total cartridge capacity for the library. **For the SL8500 and SL3000 libraries**, you must first vary the library online (Vary ACS command) to see the actual capacity of the library rather than the maximum capacity of the library. Before you vary the library, the maximum capacity is based on the highest possible number of panels that the HSC allows to be defined, not the number you defined in LIBGEN.

After you vary the library online, enter the `Display Acs` or `Display Lsm` command to show the actual library capacity. Refer to *ELS Command, Control Statement, and Utility Reference* for information about the Vary and Display commands.

Planning for the SL8500

Configuring and managing an SL8500 Library is significantly different than the other supported LSMs. Therefore, the SL8500 warrants some separate planning considerations, which are shown in the example SL8500 configuration in Step 8 through Step 10 of ["Coding the LIBGEN Macros."](#)

- The SL8500 library contains four rails on which four handbots travel, and HSC sees each rail in an SL8500 as an entire LSM. When you configure an SL8500, the SLIACS macro LSM parameter specifies the assembler labels of the 4 SLILSM macros for each SL8500 rail, where the first label corresponds to the first rail, the second label to the second rail, and so forth.

When you are defining the adjacent LSMs in an SL8500 with only internal passthru ports (elevators), the adjacent LSMs are the other rails, as shown in Step 9 in ["Coding the LIBGEN Macros."](#)

- Each SL8500 has 3 "internal PTPs" (elevators) and can also have actual PTPs that connect two SL8500s. On the SLILSM `PASSTHRU` parameter, 0 denotes an internal PTP, 1 an external PTP. The first LSM (rail) in an SL8500 is always the PTP master.
- For an SL8500 with no external PTPs, the adjacent LSMs are the other rails, as specified on the SLILSM `ADJACNT` parameter.
- HSC sees SL8500 panels as follows:
 - Panel 0 = CAP panel
 - Panel 1 = drive panel (Note: this is the only drive panel)
 - Panels 2 through n = cell storage panels. The total number of panels depends on the configuration of the library.
 - * base library — 2-10
 - * With one expansion module — 2-18 (expansion module is 8-15)
 - * With two expansion modules — 2-26 (expansion modules are 8-23)
 - * With three expansion modules — 2-34 (expansion modules are 8-31).

When you are configuring SL8500 panels, this translates into the following:

- Specify SLILSM `DRIVE=(1)` for the single drive panel.

- Specify DOOR=SL8500-1 for a single SL8500 CAP, specify DOOR=SL8500-2 for two SL8500 CAPs
- To HSC, as viewed from inside the library, SL8500 column numbers are positive starting with +1 to the right of the center line of the drive bays. They are negative starting with -1 to the left of the drive bays. HSC reports two columns for each cell storage panel (columns 0 and 1).
- storage panel (columns 0 and 1).
- To HSC, the SL8500 outer wall = 1 and the inner wall = 2.
- SL8500 row numbers: within each LSM (rail), rows are numbered consecutively from the top down. Row numbers start with 1 for the SL8500 and 0 for HSC.

Note: The SL8500 library uses TCP/IP protocol over an Ethernet physical interface to communicate with the host and HSC. Note the following:

- When SL8500s are connected through passthru ports, all hosts must communicate with only one SL8500; preferably the first or rightmost one.
 - The SL8500 library should be attached to a separate subnet or a controlled network to protect it from floods of ARP broadcasts.
-
-

Planning for Library-Attached Transports

Determine 4-digit hexadecimal MVS unit addresses for your system's library-attached transports. You will use these addresses to:

- Add a SLIDRIVS macro (ADDRESS parameter) to define RTD device addresses during the HSC LIBGEN process as described in "[Coding the LIBGEN Macros.](#)"
- Use the HCD facility to assign MVS device numbers to these transports as described in "[Configuring MVS Device Numbers and Esoterics.](#)"

Note:

- **LTO and SDLT transports are not supported as direct-attach devices in an MVS environment.** These transports are recognized by HSC but are accessible only to open systems clients using LibraryStation.
 - Nonexistent devices can be defined in LIBGEN SLIDRIVS macro only if the LIBGEN and LSM panel types match. Changing an LSM panel type can only be done by a qualified service provider. Nonexistent drives are treated by HSC as incompatible with any media type.
 - When the new devices are actually installed, be sure to reconfigure your MVS unit address and esoteric definitions to accurately reflect the changes.
 - StorageTek recommends specifying unit addresses for devices not defined on a host **if that host may be used as a server for another host where device addresses have been configured.**
 - After you define library-attached transports to HSC, you can define them as RTDs as described in ["Planning for Library-Attached Transports."](#)
-

- Refer to *Configuring and Managing SMC* for defining different unit addresses for client server.

Planning for Library-attached Transports as RTDs

After you define library-attached transports to HSC as described in ["Planning for Library-Attached Transports,"](#) you can define them as RTDs as follows:

- Specify the MVS device numbers on the CONFIG VTSS RTD DEVNO parameter.
You also specify the RTD identifier on the CONFIG VTSS RTD NAME parameter. To help identify the RTDs connected to each VTSS, StorageTek recommends that you choose RTD identifiers that reflect the VTSS name (specified on the VTSS NAME parameter) and the RTD's MVS device number (specified on the RTD DEVNO parameter).

In configurations where multiple VTSSs are connected to and dynamically share the same RTD, in each VTSS definition you can either assign unique RTD identifiers or use the same RTD identifier.
- Oracle **strongly recommends** that you define your RTDs to MVS through HCD (as normal 3490 tape drives), even if you do not intend to vary them online to MVS. This prevents the RTD addresses used in CONFIG and LIBGEN from accidentally being used for other devices. If you do not do this, and subsequently use the addresses for other MVS devices, you will cause problems with LOGREC processing, because VTCS will write records using the RTD addresses, and MVS will write records for other devices with those same addresses.
- You can specify that library-attached transports can only be used as RTDs. For more information, see ["Creating the HSC CDS."](#)
- Ensure that you use the drive operator's panel or T10000 Virtual Operator Panel (VOP) to enable the SL PROT (Standard Label Protect) function on the RTDs (ESCON or FICON).

Planning for Library Volumes

Determine the volume serial ranges and media types for VTVs and for all cartridges that will reside in the ACS. You will use these VOLSER ranges to create POOLPARM/VOLPARM statements to define real, virtual, scratch, MVC, and cleaning tape pools.

Guidelines for MVCs

- Create separate volsr range for MVCs to prevent HSC from writing to MVCs and to prevent VSM from writing to conventional Library Volumes. If you use POOLPARM/VOLPARM, VOLSER range validation is automatic.
- VTCS, not MVS, controls access to MVCs. The tape management system does not control VSM access to an MVC volume and does not record its usage. If you choose to define MVCs to the tape management system, to ensure that the tape management system does not accidentally access MVCs, follow the guidelines in ["Updating the Tape Management System for VSM."](#)
- Use your security system to restrict access to MVCs as described in ["Defining MVC Pool Volser Authority."](#)
- HSC automatically marks newly entered MVC volumes as non-scratch. If you define existing Library Volumes as MVCs, ensure that these volumes do not contain data you need, then run the HSC UNSCratch Utility to unscratch them. Refer to *ELS Command, Control Statement, and Utility Reference* for more information.
- You define MVC media and recording technique as described in ["VTCS Considerations to Correctly Specify MVC Media."](#) Note that VTCS requires a **unique value** for the MEDIA parameter on the STORCLAS statement.

VTCS Considerations to Correctly Specify MVC Media

[Table 1-5](#) describes the values required to specify the desired media and recording technique on the HSC VOLPARM/POOLARM statements and HSC STORCLAS statement to correctly specify the desired MVC media.

Table 1-5 RTD Model/MVC Media Values

Transport Model	TAPEREQ/ VOLPARM MEDIA	RECTECH	STORCLAS MEDIA	Cartridge Type - Specified by STORCLAS MEDIA	Density	Encrypted?
4490	STANDARD	STANDARD	STANDARD	standard length 3480 cartridge	single	NA
9490, 9490EE	ECART	ECART	ECART	3490E cartridge	single	NA
9490EE	ZCART	ZCART	ZCART	3490EE cartridge	single	NA
9840	STK1R	STK1RA	STK1RAB	T9840A or T9840B cartridge	single	NA
T9840B	STK1R	STK1RB	STK1RAB	T9840A or T9840B cartridge	single	NA

Table 1–5 (Cont.) RTD Model/MVC Media Values

Transport Model	TAPEREQ/ VOLPARM MEDIA	RECTECH	STORCLAS MEDIA	Cartridge Type - Specified by STORCLAS MEDIA	Density	Encrypted?
T9840C	STK1R	STK1RC	STK1RC	T9840C cartridge	double	NA
T9840D - T9840D Non- Encrypting Transport	STK1R	STK1RD	STK1RD	T9840D cartridge	triple	no
T9840DE - T9840D Encrypting Transport	STK1R	STK1RDE	STK1RDE	T9840D cartridge for encryption	triple	yes
T9940A	STK2P	STK2PA	STK2PA	T9940A cartridge	single	NA
T9940B	STK2P	STK2PB	STK2PB	T9940B cartridge	double	NA
T1A34 - T10000A Non-Encrypting Transport	T10000T1	T1A34	T1A000T1	T10000 full capacity cartridge	single	no
T1AE34 - T10000A Encrypting Transport	T10000T1	T1AE34	T1A000E1	T10000 full capacity cartridge for encryption	single	yes
T1A34 - T10000A Non-Encrypting Transport	T10000TS	T1A34	T1A000TS	T10000 sport cartridge	single	no
T1AE34 - T10000A Encrypting Transport	T10000TS	T1AE34	T1A000ES	T10000 sport cartridge for encryption	single	yes
T1B34 - T10000B Non-Encrypting Transport	T10000T1	T1B34	T1B000T	T10000 full capacity cartridge	double	no
T1BE34 - T10000B Encrypting Transport	T10000T1	T1BE34	T1B000E1	T10000 full capacity cartridge for encryption	double	yes
T1B34 - T10000B Non-Encrypting Transport	T10000TS	T1B34	T1B000TS	T10000 sport cartridge	double	no

Table 1–5 (Cont.) RTD Model/MVC Media Values

Transport Model	TAPEREQ/ VOLPARM MEDIA	RECTECH	STORCLAS MEDIA	Cartridge Type - Specified by STORCLAS MEDIA	Density	Encrypted?
T1BE34 - T10000B Encrypting Transport	T10000TS	T1BE34	T1B000ES	T10000 sport cartridge for encryption	double	yes
T1C34 - T10000C Non-Encrypting Transport	T10000T2	T1C34	T1C000T2	T10000C full capacity cartridge	triple	no
T1CE34 - T10000C Encrypting Transport	T10000T2	T1CE34	T1C000E2	T10000C full capacity cartridge for encryption	triple	yes
T1C34 - T10000C Non-Encrypting Transport	T10000TT	T1C34	T1C000TT	T10000C sport cartridge	triple	no
T1CE34 - T10000C Encrypting Transport	T10000TT	T1CE34	T1C000ET	T10000C sport cartridge for encryption	triple	yes
T1D34 - T10000D Non-Encrypting Transport	T10000T2	T1D34	T1D000T2	T10000D full capacity cartridge	quadruple	no
T1DE34 - T10000D Encrypting Transport	T10000T2	T1DE34	T1D000E2	T10000D full capacity cartridge for encryption	quadruple	yes
T1D34 - T10000D Non-Encrypting Transport	T10000TT	T1D34	T1D000TT	T10000D sport cartridge	quadruple	no
T1DE34 - T10000D Encrypting Transport	T10000TT	T1DE34	T1D000ET	T10000D sport cartridge for encryption	quadruple	yes

Use [Table 1–5](#) as a guideline to:

- Create VOLPARM/POOLARM statements that segregate single/double density media or encrypted/non-encrypted media.

- Specify the correct STORCLAS MEDIA values to assign the desired cartridge type and recording technique to MVCs.
- Determine which transport models can write to/read from which media. A higher capability transport (double density vs. single, or encryption vs. non-encryption) can read from media written by a lower capability transport but can only write to that media from the beginning of the tape. A lower capability transport, however, cannot read from media written by a higher capability transport but can write to that media from the beginning of the tape.

Examples

If you are adding T1000D encrypting transports and new media to encrypt, create new POOLPARAM/VOLPARAM statements for the new media and STORCLAS statements to allow VTCS to select this media. For example:

```
POOLPARAM NAME(SYS1MVC10D) TYPE(MVC) MVCFREE(40) MAXMVC(4) THRESH(60) INP Thrsh(10)
START(70)
VOLPARAM VOLSER(MVC900-MVC999) MEDIA(T10000T2) RECTECH(T1DE34)
STORCLAS NAME(10DENCRYPT) INPLACE(YES) MEDIA(T1D000E2)
```

Note: In the preceding example:

- The T1000D (full) MVCs have been partitioned, which allows them to be *dynamically reclaimed*.
 - Dynamic reclaim is enabled by specifying INPLACE(YES) for Storage Class 10DENCRYPT.
-

If you are adding T1000D encrypting transports and want to convert existing media to encryption media, change existing VOLPARAMs to specify encryption and change existing STORCLAS statements to request encryption. For example:

```
POOLPARAM NAME(SYS1MVC10D) TYPE(MVC) MVCFREE(40) MAXMVC(4) THRESH(60) INP Thrsh(10)
START(70)
VOLPARAM VOLSER(MVC900-MVC999) MEDIA(T10000T2) RECTECH(T1DE34)
STORCLAS NAME(T10K) INPLACE(YES) MEDIA(T1D000E2)
```

Here is how it works: If you have MVCs that already contain data, you cannot add "encrypted" VTVs to these MVCs. You can, however, encrypt data on initialized MVCs that do not contain data. To make this strategy work, therefore, ensure that you have sufficient free T1000 MVCs and also consider doing demand drains on MVCs that contain data to free them up.

Using the STORclas MEDIA Parameter for MVC Media Preferencing

By default, in mixed-media VSM systems, VTV automatic and demand migrations (and consolidations) attempt to go to MVCs by media type in this order:

1. Standard length 3480 cartridge
2. 3490E cartridge
3. 3490EE cartridge
4. T9840A/B cartridge
5. T9840C cartridge
6. T9840D cartridge

7. T9940A cartridge
8. T9940B cartridge
9. T10000A sport cartridge
10. T10000B sport cartridge
11. T10000C sport cartridge
12. T10000D sport cartridge
13. T10000A full capacity cartridge
14. T10000B full capacity cartridge
15. T10000C full capacity cartridge
16. T10000D full capacity cartridge
17. LTO-6 full capacity cartridge
18. LTO-7 full capacity cartridge
19. LTO-8 full capacity cartridge

By default, for automatic and demand space reclamations, VSM attempts to write VTVs to output MVCs by media type in this order:

1. LTO-8 full capacity cartridge
2. LTO-7 full capacity cartridge
3. LTO-6 full capacity cartridge
4. T10000D full capacity cartridge
5. T10000C full capacity cartridge
6. T10000B full capacity cartridge
7. T10000A full capacity cartridge
8. T10000D sport cartridge
9. T10000C sport cartridge
10. T10000B sport cartridge
11. T10000A sport cartridge
12. T9940B cartridge
13. T9940A cartridge
14. T9840D cartridge
15. T9840C cartridge
16. T9840A/B cartridge
17. 3490EE cartridge
18. 3490E cartridge
19. Standard length 3480 cartridge

The MEDIA parameter of the STORclas statement specifies a preference list of MVC media types. This list supersedes the default media selection list. Note that for reclamation, VTCS attempts to write VTVs back to MVCs in the reverse of the order specified on the MEDIA parameter. For more information about STORCLAS MEDIA values, see ["VTCS Considerations to Correctly Specify MVC Media."](#)

For example, if you specify the following on the MEDIA parameter of the STORclas statement:

```
MEDIA (STK1RAB, STK1RC, STK2PB)
```

- To select an MVC for migration to this Storage Class, VTCS searches for a usable MVC in the order STK1RAB, STK1RC, STK2PB.
- To select an MVC for the output of reclaim to this Storage Class, VTCS searches for a usable MVC in the order STK2PB, STK1RC, STK1RAB.

You can specify the media and ACS preferencing through the Storage Class(es) specified on the MIGpol parameter of the MGMTclas control statement.

To optimize recall processing in mixed-media systems, ensure that your MVC pool has at least one media type compatible with each RTD type.

Planning for VTSSs

Determine your system's 1 to 8 character VTSS IDs, which you specify when you run VTCS CONFIG to install and configure your VSM system as described in "[Configuring VTCS](#)."

Caution: Note the following:

- The VTSS ID can consist of the characters “A-Z”, “0-9”, “@”, “\$”, and “#”.
- You specify the VTSS ID **only** through the NAME parameter, which sets the VTSS ID in both the VTSS microcode (as displayed in the Subsystem Name field in the LOP or VOP) and in the configuration area of the HSC CDS. After VSM is put into operation, the VTSS ID is also stored in each VTV record in the CDS. Each VTV record contains either:

The VTSS ID on which that VTV is resident.

The VTSS ID from which the VTV was migrated.

- Once you set the VTSS ID using the NAME parameter, you **cannot** change this identifier in the HSC CDS. That is, the CONFIG utility **does not** let you change the NAME parameter after an initial setting. Moreover, changing the VTSS ID using the Subsystem Name field of the LOP, VOP, or DOP **cannot** change the VTSS ID in the HSC CDS.
 - It is especially critical that you **do not** attempt to rename a VTSS that contains data on VTVs, which includes VTSS-resident VTVs and migrated VTVs.
 - For an initial setting **only** (not a change), you can set the VTSS ID in the NAME parameter only if the VTSS ID value in the VTSS microcode is:
 - The factory setting (all blanks).
 - A value of 99999999 (eight 9s).
 - Therefore, for an initial setting **only**, if the name in the VTSS microcode **is not** all blanks or 99999999, your StorageTek hardware representative must use the VTSS LOP, VOP, or DOP to set the VTSS ID to 99999999. This will enable you to set the VTSS ID to the value you want using the NAME parameter.
-
-

Planning for VTDs

Determine MVS unit addresses for your system’s VTDs as follows:

- For each VTSS in your VSM configuration, determine a unique unit address range for the VTDs in that VTSS. Do not use duplicate addresses or overlapping address ranges, either within the VTDs in a VTSS or across VTSSs.
- For each VTSS in your VSM configuration, define its VTD unit addresses to VTCS using CONFIG.

In a multi-host, multi-VTSS configuration, you can configure your VTD unit addresses to restrict host access to VTSSs. Note that the VTVs created and MVCs initially written to from a VTSS are considered that VTSS’s resources, so only hosts with access to a VTSS also have access to its VTVs and MVCs. For more information, see "[Configuring VTCS](#)." However, restricting access by host restricts your ability to use the client/server feature. Also, if you are using client/server, you will need to define host access to VTDs that do not have data path access to some drives (as these drives will be used by client hosts). In that case, you need to use the NOVERIFY parameter when the VTDs are defined to server hosts.

- For each SMC host, use the HCD facility to define to MVS the VTDs that host can access as described in "[Configuring MVS Device Numbers and Esoterics](#)." The unit addresses you specify through the HCD facility **must** match the unit address range you specified for that host using CONFIG, or use the SMC unit address mapping facility to use different addresses. Refer to *Configuring and Managing SMC* for more details.
- If you use CA-MIA or MVS shared tape feature, add VTDs to the list of managed devices.

Planning for VTVs

Use these guidelines to plan for VTVs:

- Depending on the VTV size you choose and the media your volumes were previously written on, you may need to update the JCL volume count. You can do this without changing the JCL. For more information, see the IDAXVOLCNT parameter in *Configuring and Managing SMC*.
- HSC/VTCS **does not** allow allocation of unlabeled tapes to VTVs. Unlabeled VTVs can cause the following for scratch VTV allocation requests:
 - If your JCL specifies a virtual esoteric, SMC fails the allocation.
 - If you have a default esoteric such as CART and specify allocation to virtual (using SMC policy), the allocation will go to a non-virtual device.
- You must define VTV volsers to your tape management system; for more information, see "[Updating the Tape Management System for VSM](#)."
- Ensure that VTV volsers ranges do not duplicate or overlap existing TMS ranges or volsers of real tape volumes, including Library Volumes, **including** MVCs and Library Volumes that are regularly entered and ejected from the ACS.
- If VTDs are being used across multiple MVS images and VTV volsers are unique, add a generic entry for SYSZVOLS to the SYSTEM inclusion RNL to insure that a VTV is used by only one job at a time. If you are using automatic tape switching, also add a generic entry for SYSZVOLS to the SYSTEM inclusion RNL to prevent a system from holding a tape device while it waits for a mount of a volume that is being used by another system.

For more information, refer to the IBM publication *z/OS MVS Planning: Global Resource Serialization*.

- HSC support lets you mix VTV and real volume types in the same scratch pool. In this case, ensure that the mount request specifies a VTD as transport type (for example, using POLICY MEDIA (VIRTUAL)). In addition, if you are routing data to a specific VTSS (for example, by using esoteric substitution) and the request specifies a subpool, ensure that the subpool contains scratch VTVs.

Refer to *Managing HSC and VTCS* for information about displaying and managing scratch subpools.

- By default, VTCS assigns a Management Class to VTVs only on scratch mounts. You can, however, specify that VTCS assigns a Management Class whenever VTCS mounts a VTV (for read or write).

Caution: If you specify that VTCS assigns a Management Class whenever VTCS mounts a VTV, these attributes can change, which can cause undesirable or unpredictable results.

For example, if an application writes data set PROD.DATA to VTV100 with a Management Class of PROD, then writes data set TEST.DATA to VTV100 with a Management Class of TEST, then the VTV (and both data sets) has a Management Class of TEST. Similarly, it is possible to write SMC POLICY statements that assign different management classes to the same data set (for example, based on jobname), which can also cause a VTV's Management Class to change.

- To accommodate your VSM system's VTVs, you may need to increase the DASD space for your tape management system. After you determine the number and range of VTVs your VSM system requires, refer to your tape management system documentation for specific information on calculating the DASD space requirements.

Data Chaining a VTD Read Forward or Write Command

Note that when data chaining a Read Forward or Write command, the VTSS requires the minimum data chained update count.

Planning VTCS Operating Policies

This chapter describes the VTCS operating policies. You must set your VTSS AMTs, as described in "AMT Settings." Otherwise, you can just take the defaults, let the system run for a while, and see what is working or not working and adjust accordingly.

The following sections describe VTCS operating policies:

- "Making CDS Logging Required"
- "Locked Resource Policy"
- "VTSS Policies"
- "General VTV Policies"
- "VTV Migration Policies"
- "VTV Replication Policies"
- "MVC Policies"

Note that these are **rough groupings** only. All the pieces and parts of VSM interact with each other constantly, so it is impossible to talk VTVs without thinking about MVCs, and vice versa. These policies are set globally on the CONFIG statement¹, some of which you can override on a per job basis.

Making CDS Logging Required

Table 2-1 Enable CDS Logging

Policy:	Values:	Default:	Command to Set Policy:
Makes CDS logging required.	<ul style="list-style-type: none"> ▪ LOGPOL=REQuired ▪ LOGPOL=OPTional 	LOGPOL=OPTional	CONFIG GLOBAL LOGPOL parameter

Usage Notes

- If all copies of the CDS fail, recovery can be time consuming. CDS logging can aid in re-synchronizing the CDS with the data content of the VSM subsystem. Note that recovery using logs will recreate a usable CDS but it **may not** fully restore VTV and MVC metadata.
- Note that the CONFIG GLOBAL LOGPOL setting merely specifies whether logging is required or optional. For information on setting up logging, see "Defining and Creating CDS Log Files." Refer to *Managing HSC and VTCS* for more

¹ There is one exception, VTV size (MAXVtvsz on MGMTclas).

information. For your initial configuration, you may want to go with the default (LOGPOL=OPT) and change this value, if desired, when you actually decide to implement logging.

- Logging works as follows:
 - If you specify LOGPOL=REQ and all hosts sharing the CDS are at 7.0 or above, logging is enabled for all systems, which aids recreation of a CDS.
 - If you specify LOGPOL=REQ and **some** hosts sharing the CDS are at below 7.0, logging is enabled for **only** 7.0 and above systems, so recovery is from logs for **only** those hosts. On hosts below 7.0, VTCS does not initialize.
 - If you specify LOGPOL=OPT and **some** hosts sharing the CDS are at below 7.0, logging is enabled for **only** 7.0 and above systems, so recovery is from logs for **only** those hosts. Hosts below 7.0 initialize but do not do logging.

Running IMPORT/EXPORT When Logging Required is Enabled

To run IMPORT/EXPORT when Logging Required is enabled:

1. Back up the CDS.
2. Run CONFIG to turn off logging required.
3. Run the utility.
4. Run CONFIG to re-enable logging required.
5. Back up the CDS.
6. Resume normal activity.

Locked Resource Policy

Table 2–2 shows the locked resource policy.

Table 2–2 *Locked Resource Policy*

Policy:	Values:	Default:	Command to Set Policy:
Specifies the minimum number of minutes that a resource is locked before message SLS6946E is issued.	Valid values are 0, or any value between 5 and 240. If 0 is specified, message SLS6946E will not be issued when a required resource is locked.	If this parameter is not specified, the current default of 10 minutes is retained.	CONFIG LOCKTOUT= <i>nnn</i>

Note: LOCKTOUT is only supported at 'F' level CDS (V61ABOVE) and above.

VTSS Policies

The following sections describe these VTSS policies:

- ["AMT Settings."](#) AMT settings are critical because they start and stop automatic migration and all its attendant functions. On this one, you have CONFIG settings that you can override with the SET MIGOPT command.

- **"Deleting Scratched VTVs."** Deleting scratched VTVs can free buffer space but has some associated potential risks, so read the usage notes before enabling this feature.
- **"VTV Page Size" and "Maximum VTV Size."** These two are linked, and even though they are VTV policies, they are more about optimizing VTSS usage, so they are discussed here. Large VTV page sizes, when matched with the appropriate VTSS model, can optimize performance within the VTSS. VTV page sizes are about the page data size within the VTSS, not the actual VTV size. VTV page sizes, however, are actively linked to VTV sizes, which can improve performance in high-capacity VTSSs. In fact, selecting VTV sizes of 2 or 4 GB enforces a Large VTV Page.
- **"Maximum RTDs per VTSS."** Your choices are 16 or 32 maximum, and it is a CONFIG GLOBAL setting.

AMT Settings

Table 2–3 shows AMT settings.

Table 2–3 AMT Settings

Policy:	Values:	Default:	Command to Set Policy:
Controls the automatic space management/migration cycle. This cycle begins when DBU (Disk Buffer Utilization) exceeds the high AMT (HAMT) or as described in Table 2–4. AMTs are a percentage of DBU.	<ul style="list-style-type: none"> ■ LAMT - 5 - 95, must be one less than HAMT. ■ HAMT - 5 - 95, must be one greater than LAMT. 	<ul style="list-style-type: none"> ■ LAMT - 70 ■ HAMT - 80 	<ul style="list-style-type: none"> ■ CONFIG VTSS LOW and HIGH parameters ■ SET MIGOPT HIGHthld and LOWthld parameters.

Table 2–4 shows VTSS Model Maximum VTVs and Automigration Thresholds.

Table 2–4 VTSS Model Maximum VTVs and Automigration Thresholds

VSM Model	Maximum VTVs	Automigration Starts Above
VSM2/VSM3	100,000	97,000
VSM4/VSM5	300,000	291,000
VSMc Models:	500,000	485,000
<ul style="list-style-type: none"> ■ VSM5-45TB-IFF3 ■ VSM5-68TB-IFF3 ■ VSM5-90TB-IFF3 		
VSM 6	1,000,000	970,000

Usage Notes

- With CONFIG, AMT settings take effect when you start HSC and apply to the specified VTSS.
- With SET MIGOPT:

- AMT settings take effect immediately and apply to the specified VTSS or if no VTSS is specified, to all VTSSs. If you try to set global values (no VTSS specified) and the values are not valid for one VTSS (for example, MAXMIG(5) and one VTSS only has 4 RTDs connected), VTCS will not set values for any VTSSs.
- You can set the LAMT, the HAMT, or both.
- The following are general guidelines for changing the defaults:
 - The difference between the high and low AMTs affects the duration of the space management/migration cycle.
 - Lowering the HAMT tends to trigger more frequent space for management/migration cycles.
 - Raising the HAMT tends to trigger less frequent space management/migration cycles.
 - Lowering the LAMT tends to free more VTSS space and migrate more VTVs.
 - Raising the LAMT tends to keep more VTVs resident in VTSS space and migrate fewer VTVs.

Tip: You can use Display VTSS to display the DBU, HAMT, and LAMT for each VTSS in your system. You can also use Display MIGrate to display migration status.

Deleting Scratched VTVs

Table 2-5 Delete Scratched VTVs Settings

Policy:	Values:	Default:	Command to Set Policy:
Specifies whether VSM deletes scratched VTVs.	NO, YES, FORCE	NO	MGMTCLAS DELSCR

Usage Notes

- Specifying DELSCR(YES) deletes scratch VTVs unless this would cause a processing delay due to unavailable resources.
- Specifying DELSCR(FORCE) waits for access to any resources that are required and deletes scratch VTVs, which may cause a processing delay.
- Specifying DELSCR(YES) or DELSCR(FORCE) causes VSM to delete scratched VTVs, which frees VTSS buffer space and (logically) deletes any VTV copies from MVCs so that MVC space can be reclaimed.

Note: As an alternative, you can take the default DELSCR(NO) and, as needed, run the DELETSCR NOTREF command to delete scratch VTVs (from the VTSS and MVCs) n days after the VTV was last referenced. Refer to *ELS Command, Control Statement, and Utility Reference*. for more information.

Caution:

- When you scratch a VTV with DELSCR YES or FORCE attribute, VSM erases the VTV data at scratch synchronization time, which eliminates the ability to "unscratch" a VTV to recover data.
- When using SLUCONDB to synchronize scratches, it is important to avoid the following to avoid inadvertently scratching a VTV resulting in data loss at scratch synchronization time:
 - Running the HSC SLUADMIN Scratch Update Utility at the same time that SLUCONDB is running.
 - Not specifying the current TMS database and/or the current HSC CDS when using SLUCONDB.

For more information about HSC scratch synchronization with the Scratch Conversion Utility (SLUCONDB), refer to the *ELS Command, Control Statement, and Utility Reference*.

Also note that for HSC and MVS/CSC, the DELDISP parameter has two values that affect how HSC manages the scratch status of VTVs and real volumes that were mounted scratch and the delete disposition on the dismount message is delete ('D').

- For more information about LCM scratch synchronization with the SYNCVTV function, refer to the *LCM User's Guide*.

VTV Page Size

Table 2–6 VTV Page Size

Policy:	Values:	Default:	Command to Set Policy:
Specifies the default page size used to store VTV data in the VTSS and on the MVCs for 400 and 800 MB VTVs only. For 2 GB and above VTVs, a VTVPAGE setting of LARGE is always used.	STANDARD or LARGE	STANDARD	CONFIG GLOBAL VTVPAGE MGMTclas VTVPAGE

Usage Notes

- Large page size (which requires that the CDS is at a G level or above) can provide improved performance within the VTSS and for migrates and recalls. For more information about CDS levels, see "[CDS VTCS Level Requirements](#)."

Caution:

- The page size of a VTV can **only** be changed by a VTV scratch mount. Additional restrictions may also apply for scratch VTVs that were previously resident in a VTSS.
- VTVPAGE **does not** apply to VSM2s. The VTSS microcode requirements are as follows:
 For VSM3s: microcode level N01.00.77.00 or higher.
 For VSM4s/VSM5s: microcode level D02.02.00.00 or higher.
 For VSM6s, all levels of microcode.
- If you specify LARGE and the CDS level and/or VTSS microcode **do not** support LARGE, VTCS issues warning messages and VTVPAGE defaults to STANDARD.
- Creating VTVs with large pages makes these VTVs **unreadable** in configurations that do not support large VTV pages.

Maximum VTV Size

Table 2-7 Maximum VTV Size

Policy:	Values:	Default:	Command to Set Policy:
Specifies the maximum size of a VTV.	400 -	400	CONFIG GLOBAL MAXVTVSZ,
	400 Mb.		
	800 -	800 Mb. The CDS must be at a D level or above.	MGMTCLAS MAXVTVSZ
	800 Mb. The CDS must be at a D level or above.		
	2000 -	2 Gb. The CDS must be at a G level or above.	
	2 Gb. The CDS must be at a G level or above.		
	4000 -	4 Gb. The CDS must be at a G level or above.	
	4 Gb. The CDS must be at a G level or above.		
32000 -	32GB. The CDS must be at a I level or above and this only applies to VSM 6.		
32GB. The CDS must be at a I level or above and this only applies to VSM 6.			

Usage Notes

- The size of a VTV changes only after it goes through a scratch cycle. Therefore, if you change the Management Class and DISP=MOD, then it will still retain the original size.
- If you specify a VTV size that is not supported by the configuration, VTCS issues warning messages and MAXVtvsz defaults to the largest VTV size supported by the configuration.

- MAXVtvsz **does not** apply to VSM2s.

Caution: Specifying 2GB, 4GB, or 32GB VTVs:

- Increases MVC usage.
 - May degrade space management of smaller-capacity VTSSs, which have relatively small cache and buffer sizes.
 - Increases delays to jobs waiting recalls. Although the actual data transfer time from an MVC is the same and there may be fewer interruptions, each interruption lasts longer, which can cause job time-outs.
-
- The CONFIG GLOBAL and MGMTCLAS MAXVTVSZ parameters interact as follows:
 - If MAXVTVSZ is specified on MGMTCLAS, this value overrides the CONFIG GLOBAL MAXVTVSZ value.
 - If MAXVTVSZ is **not** specified on MGMTCLAS, the CONFIG GLOBAL MAXVTVSZ value, if specified, is used. Otherwise, MAXVTVSZ defaults to 400MB.
 - If MAXVTVSZ is **not** specified on MGMTCLAS **or** on CONFIG GLOBAL, MAXVTVSZ defaults to 400MB.

Maximum RTDs per VTSS

You can connect up to 32 RTDs per VTSS.

Table 2–8 Maximum RTDS per VTSS - 16 or 32

Policy:	Values:	Default:	Command to Set Policy:
Specifies the maximum RTDs per VTSS.	16, 32	16	CONFIG GLOBAL MAXRTDs

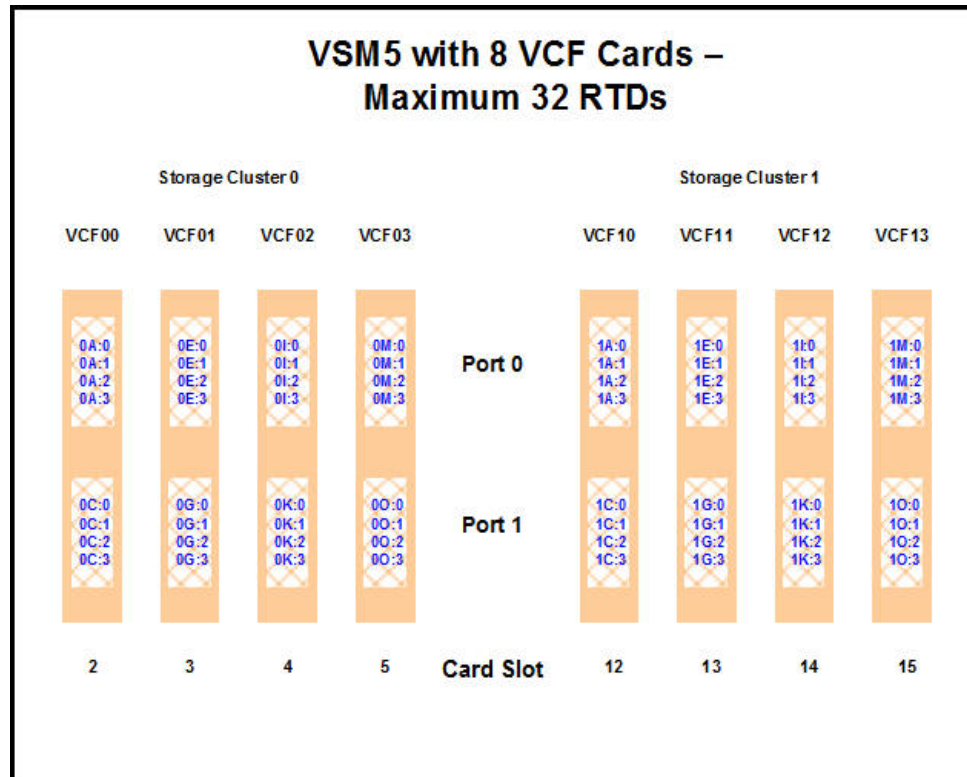
Usage Notes

- For maximum 32 RTDs support, you must fulfill the requirements described in [Table H–3](#).
- The VTCS addressing scheme for maximum 32 RTDs is different than that for maximum 16 RTDs. For more information, see "[RTD/CLINK Addresses - Maximum 32 RTDs](#)."

RTD/CLINK Addresses - Maximum 32 RTDs

VSM5 is available only with 8 VCF (FICON) cards in the configuration for a maximum of 32 RTDs shown in [Figure 2–1](#).

Figure 2–1 VSM5 with 8 VCF cards - Max 32 RTDs



As Figure 2–1 shows, each FICON interface supports 4 devices attached using FICON directors. Note that the addressing scheme used by VTCS is different from that used for support of 2 devices per interface. The device addresses are now in the format *CI:R*, where:

- C is the cluster number (0 or 1)
- I is the interface number (A, C, E, G, I, K, M, or O)
- R is the device number on the interface (0, 1, 2, or 3).

Table 2–9 shows the "old" address (maximum 16 RTDs) and its corresponding "new" address (maximum 32 RTDs).

Note: If you upgrade an existing configuration from maximum 16 to maximum 32 RTDs, you have to change the addresses in your CONFIG deck on your RTD statements, your CLINK statements, or both.

Table 2–9 RTD/CLINK Addresses - Maximum 32 RTDs

Cluster Number	Interface	RTD/CLINK	Old Address (Maximum 16 RTDs)	New Address (Maximum 32 RTDs)
0	A	0	0A	0A:0
0	A	1	0B	0A:1
0	A	2	-	0A:2

Table 2–9 (Cont.) RTD/CLINK Addresses - Maximum 32 RTDs

Cluster Number	Interface	RTD/CLINK	Old Address (Maximum 16 RTDs)	New Address (Maximum 32 RTDs)
0	A	3	-	0A:3
0	C	0	0C	0C:0
0	C	1	0D	0C:1
0	C	2	-	0C:2
0	C	3	-	0C:3
0	E	0	0E	0E:0
0	E	1	0F	0E:1
0	E	2	-	0E:2
0	E	3	-	0E:3
0	G	0	0G	0G:0
0	G	1	0H	0G:1
0	G	2	-	0G:2
0	G	3	-	0G:3
0	I	0	0I	0I:0
0	I	1	0J	0I:1
0	I	2	-	0I:2
0	I	3	-	0I:3
0	K	0	0K	0K:0
0	K	1	0L	0K:1
0	K	2	-	0K:2
0	K	3	-	0K:3
0	M	0	0M	0M:0
0	M	1	0N	0M:1
0	M	2	-	0M:2
0	M	3	-	0M:3
0	O	0	0O	0O:0
0	O	1	0P	0O:1
0	O	2	-	0O:2
0	O	3	-	0O:3
1	A	0	1A	1A:0
1	A	1	1B	1A:1
1	A	2	-	1A:2
1	A	3	-	1A:3
1	C	0	1C	1C:0
1	C	1	1D	1C:1
1	C	2	-	1C:2

Table 2–9 (Cont.) RTD/CLINK Addresses - Maximum 32 RTDs

Cluster Number	Interface	RTD/CLINK	Old Address (Maximum 16 RTDs)	New Address (Maximum 32 RTDs)
1	C	3	-	1C:3
1	E	0	1E	1E:0
1	E	1	1F	1E:1
1	E	2	-	1E:2
1	E	3	-	1E:3
1	G	0	1G	1G:0
1	G	1	1H	1G:1
1	G	2	-	1G:2
1	G	3	-	1G:3
1	I	0	1I	1I:0
1	I	1	1J	1I:1
1	I	2	-	1I:2
1	I	3	-	1I:3
1	K	0	1K	1K:0
1	K	1	1L	1K:1
1	K	2	-	1K:2
1	K	3	-	1K:3
1	M	0	1M	1M:0
1	M	1	1N	1M:1
1	M	2	-	1M:2
1	M	3	-	1M:3
1	O	0	1O	1O:0
1	O	1	1P	1O:1
1	O	2	-	1O:2
1	O	3	-	1O:3

General VTV Policies

The following sections describe these VTV policies, roughly in order of importance:

- ["Maximum VTVs per MVC."](#) Take the default, or see the discussion about high-capacity media.
- ["Hosts Disabled from Migration, Consolidation, and Export by VTV or Management Class."](#) If you do not want a host to do these things, you can explicitly turn them off. The default is not to disable any hosts.
- ["Recall VTVs with Read Data Checks."](#) The default is to recall VTVs with Read Data Checks. If this does not work for you, you can turn it off either globally or per specific operation.

- ["Early Time to First Byte \(ETTFB\)."](#) ETTFB (also known as the concurrent tape recall/mount feature) allows host applications to read data while VTVs are being recalled from RTDs.
- ["VTV VOLSAFE Policy,"](#) which allows you to set the VTSS-resident VTV write policy.

Maximum VTVs per MVC

Table 2–10 *Maximum VTVs per MVC*

Policy:	Values:	Default:	Command to Set Policy:
Specifies the maximum number of VTVs that can be migrated to a single MVC.	<ul style="list-style-type: none"> ■ 4 to 32000 for a D, E or F level CDS. ■ 4 to 65000 for a G level or higher CDS. 	<ul style="list-style-type: none"> ■ 32000 for a D, E or F level CDS. ■ 65000 for a G level or higher CDS. 	CONFIG GLOBAL MAXVTV parameter

Usage Notes

- This policy applies to all MVCs and, at the time you set the policy, applies only to future migrations. That is, it will not lower the number of VTVs already migrated to an MVC. If the policy is not specified, the default is as shown in [Table 2–10](#) unless the available MVC space is less than any remaining current VTSS resident VTV.
- Generally, use the default to allow VSM to automatically manage VTV stacking. However, with high-capacity media (for example, in a VSM system where all MVCs are type STK2P), specifying a value lower than the maximum may improve recall performance in some situations. Note, however, that a very low value can reduce that percentage of usable MVC space. If the maximum VTVs per MVC is exceeded, then usable space is reported as 0%.
- For more information about CDS levels, see ["CDS VTCS Level Requirements."](#)

Hosts Disabled from Migration, Consolidation, and Export by VTV or Management Class

Table 2–11 *Hosts Disabled from Migration, Consolidation, or Export by VTV or MGMTclas*

Policy:	Values:	Default:	Command to Set Policy:
Specifies that a host cannot initiate automatic and demand migration and consolidation processing, or export by VTV or Management Class.	NOMIGRAT, if not specified, migration, etc., is not disabled .	Do not disable	CONFIG HOST

Usage Notes

- Specifying NOMIGRAT also causes NORECLAM to be set; for more information, see ["Hosts Disabled from Reclamation."](#)

- MGMTclas IMMEDmig KEEP and IMMEDmig DELETE are mutually exclusive with CONFIG HOST NOMIGRAT. If you specify both, the IMMEDmig value overrides NOMIGRAT (for only those VTVs with the IMMEDmig value), and VTCS does not issue a message about this override.

Recall VTVs with Read Data Checks

Table 2–12 Recall VTVs with Read Data Checks

Policy:	Values:	Default:	Command to Set Policy:
Specifies whether VTCS recalls VTVs with read data checks.	YES - Recall VTVs with Read Data Checks.	NO	CONsolid RECALWER
	NO - Do not recall VTVs with Read Data Checks.		EEXPORT RECALWER MVCDRAIN RECALWER RECALL RECALWER

Usage Notes

- The global default is RECALWER=NO, and cannot be overridden at the global level. It can be overridden by specifying RECALWER=YES at the command level as described in [Table 2–12](#).
- During MVC reclaims, VTCS will never recall VTVs with read data checks, regardless of the RECALWER setting.
- If a recall fails with a data check, and RECALWER=YES, VTCS attempts to recover the data, which includes both swapping to alternative MVC volumes (if more than one copy of the VTV was migrated) and swapping to other RTDs. If all attempts to recover the data has failed, VTCS has two options:
 - Recall the VTV with the data check. This means that when the application attempts to read the data, a data check is presented to the application and the application can take what ever recovery steps are required (for example, try to operate just like a real tape).
 - Fail the recall, in which case the VTV will not mount and the application will time out.

Early Time to First Byte (ETTFB)

Table 2–13 ETTFB Policy

Policy:	Values:	Default:	Command to Set Policy:
Specifies whether VTCS recalls VTVs with read data checks.	YES, NO	NO	CONFIG GLOBAL FASTRECL

Usage Notes

ETTFB (also known as the concurrent tape recall/mount feature) allows host applications to read data while VTVs are being recalled. ETTFB applies to recall from MVCs and from VMVCs (VLE 1.2 and above). ETTFB is done by overlapping the VTV recall and mount phases, which allows the application to read the VTV data sooner. If

the application attempts to read part of the VTV that has not been recalled, the application's I/O request is blocked until the required VTV data is been recalled.

ETTFB is a good choice for applications that serially access the VTV data. ETTFB is generally not a benefit for those applications that stack multiple files on a single VTV, including HSM and image management applications. In these types of applications, the desired data is usually not at the beginning of the VTV, but rather at some random location in the VTV

CONFIG GLOBAL FASTRECL=YES is the global setting to enable ETTFB, which you can override on a per VTSS basis using the CONFIG VTSS NOERLYMT parameter.

VTVs that have incurred an ETTFB recall error have an error flag set in their VTV record in the CDS. These VTVs are subsequently not selected for ETTFB. If you want to reset the error flag, do the following:

1. Enter a VTMMAINT SCRATCH(ON) command for the VTV.
2. Migrate the VTV to a new MVC copy.
3. Import the VTV.
4. Create a new version of the VTV.
5. Scratch the VTV.

VTV VOLSAFE Policy

This policy specifies the VTV VOLSAFE policy.

Table 2-14 Stacked Migrates Policy

Policy:	Values:	Default:	Command to Set Policy:
Specifies the VTV VOLSAFE policy.	MANY, ONCE	MANY	MGMTclas WRITE parameter

Usage Notes

The MGMTclas WRITE parameter specifies the VTSS-resident VTV VOLSAFE policy as follows:

MANY

the default, specifies no VOLSAFE write protection.

ONCE

after the VTV is non-scratch, it cannot be overwritten or appended while it is mounted on a VTD.

Note:

- Migrated VTVs no longer have VOLSAFE protection unless you migrate them to VOLSAFE MVC media.
- Scratch processing generates an explicit security call to check the user's authority when attempting to scratch a VOLSAFE protected VTV. RACF ALTER authorization must be established for any userid used to scratch VOLSAFE protected volumes.

You can now use HSC HSC Volume Access User Exit (SLSUX14) to prevent a Security Access Facility (such as RACF) from putting volumes into scratch status. Refer to the publications *ELS Programming Reference* and *Managing HSC and VTCS* for more information.

VTV Migration Policies

The following sections describe VTV Migration policies, and end with a comprehensive example in ["Putting It All Together: VTV Migration Control Example for an Enterprise."](#)

- ["Streamed/Stacked Migrates."](#) Enabling stacked migrates can improve migration performance, and it is a CONFIG GLOBAL setting.
- ["Minimum Residency Time before Automatic Migration Candidacy."](#) A method to influence VTV availability. You can also influence how quickly VTVs are automigrated after the interval expires per ["Maximum Residency Interval before Deletion."](#)
- ["Maximum Residency Interval before Deletion."](#) The opposite of the above bullet, a method to quickly safeguard your data and conserve VTSS space.
- ["Immediate Migration Delay Interval."](#) This specifies the amount of time in minutes to wait until immediate migrate selects VTVs for migration.
- ["Maximum and Minimum Concurrent Migration Tasks."](#) Something to pay attention to because a thoughtful setting can help optimize your VSM resources. Again, you have a global setting that you can override with the SET MIGOPT command.
- ["Prioritizing Automatic and Immediate Migration Workloads."](#) A carefully planned setting can help optimize your VSM resources. Higher values may speed up migration, lower values optimize cartridge use for volumes being sent offsite.
- ["Controlling Migration Workloads,"](#) which is the inverse of the above bullet, and applies to the full range of migration types.
- ["Conditionally Change Migration Preferencing/Depreferencing."](#)
- ["Dependent Immediate Migration,"](#) which you can use to control VLE to VLE copy. For more information, refer to *Configuring the Host Software for VLE*.

Tip: You can use Display MIGrate to display migration status.

Streamed/Stacked Migrates

This policy enables streamed/stacked migrates.

Table 2–15 Stacked Migrates Policy

Policy:	Values:	Default:	Command to Set Policy:
Specifies whether streamed/stacked migrates is enabled.	STREAM, STACKED NO	NO	CONFIG GLOBAL FASTMIGR

Usage Notes

Enabling the stacked migrates feature can improve migration performance as follows:

STREAM

The streaming method is used for migrations. VTCS monitors the responses from the RTD and uses them to decide as to when a VTV has become migrated. Full advantage is made of the buffer within the RTD to improve the throughput when performing migration.

This option also implies the use of the STACKED feature as well.

STACKED

The stacked method is used for migrations. VTCS maintains a small queue of requests to the VTSS. Advantage will be taken of the various buffers in the VTSS and RTD to improve the throughput when performing a migration. For backward compatibility, the value YES is also taken to mean STACKED.

No

Disable stacked migrates (the default).

Note: FASTMIGR=STREAM | STACKED has the following prerequisites:

- FASTMIGR=STACKED. VSM4/VSM5 microcode D02.05.00.00 or higher. If this level of microcode is not installed on all VTSSs in the configuration, Stacked Migration will be limited to the VTSSs that have it installed.
- FASTMIGR=STREAM. VSM4/VSM5 microcode D02.15.xx.00 or higher. If this level of microcode is not installed on all VTSSs in the configuration, Streamed Migration will be limited to the VTSSs that have it installed.
- ELS 7.0 or higher with PTFs.
- CDS level G or higher.
- FICON ports for FICON RTDs and CLINKs.

For the Stacked Migration feature to be enabled, all hosts must be running the prerequisites, otherwise:

- If a host is active that does not support or tolerate stacked migrates, this will cause the CONFIG utility to return an error.
 - If a host is started and does not tolerate or support this feature, the host will shut down.
-
-

Minimum Residency Time before Automatic Migration Candidacy

Table 2–16 Minimum VTV Residency Interval

Policy:	Values:	Default:	Command to Set Policy:
By default, VSM selects VTVs for automatic migration based on size and usage. You can, however, specify a VTV's minimum residency time before automatic migration candidacy.	1 - 9999	None	MGMTCLAS RESTIME

Usage Notes

- The RESTIME value in a VTV's Management Class sets the recommended interval that the VTV remains VTSS-resident from the time that instance of the VTV is created, which is a **method to influence VTV availability**. A new instance of the VTV is created whenever the VTV is updated. At automigration time, the creation date and time of the VTV instance plus the RESTIME value is compared to the TOD clock to determine if the VTV is an automatic migration candidate.
- Note the following:
 - A VTV's Management Class (and attributes, such as RESTIME) is set after a scratch mount or optionally after a specific mount if `VTVattr = ALLmount`.
 - The RESTIME value is only a recommendation. VTCS can migrate a VTV before its residency interval expires if the DBU has not reached the LAMT or the specified migrate-to-threshold value and no VTVs have expired their residency intervals.
 - You can do a demand migrate of a VTV and delete it from the VTSS even if its residency interval has not expired.

The following scenario shows how the RESTIME parameter works:

1. You create Management Class with a RESTIME of 10 hours.
2. A job requests a scratch mount for the Management Class you created in Step 1. VTCS selects and mounts a scratch VTV. The VTV is updated, so at dismount time, its RESTIME value is set to 10 hours (which began when VTCS mounted the VTV).
3. VTCS migrates the VTV after 3 hours, then recalls the VTV 2 hours later for a read. The RESTIME value is not reset, and there are now 5 hours of residency remaining.
4. Two hours later, a job updates the VTV, which was 7 hours old. The update creates a new instance of the VTV and the residency interval will restart from the time the VTV was mounted for update.
5. Twenty four hours later, VTCS migrates the VTV, then recalls it 2 days later for a read. VTCS does not create a new instance of the VTV because it is not updated. The residency interval has expired and the VTV is therefore an automatic migration candidate based only on least-recently-used/size criteria.
6. A week later, the VTV is scratched. VTCS eventually selects and mounts the VTV to satisfy a scratch mount request. If the VTV is updated, its residency interval is set to the RESTIME value of the Management Class being used.

Maximum Residency Interval before Deletion

Table 2–17 *Maximum Residency Interval before Deletion*

Policy:	Values:	Default:	Command to Set Policy:
Specifies the maximum residency interval for migrated VTVs before deletion from the VTSS.	0 - 9999 hours	9999	MGMTCLAS DISCARD

Usage Notes

DISCARD lets you improve VTSS buffer management by requesting that VTVs that will not be used after a known time interval are preferenced for auto migration.

DISCARD applies to automatic and immediate migration requests as follows:

- At the start of auto migration, VTVs are preferenced for deletion from the VTSS if the DISCARD value is expired and all required migrated copies exist on MVCs. If the DISCARD value is not expired, the VTV is treated the same as a VTV that does not have a DISCARD value assigned.
- For immediate migrate requests, whether a VTV is deleted is based on comparisons of the IMMDELAY and DISCARD times when the VTV is initially placed on the immediate or delayed migration queues. For example, if a class of data is created and referenced immediately and you set both IMMDELAY and DISCARD to 2 hours, the VTVs are immediately migrated and deleted from the VTSS after 2 hours.
- At the start of auto migration, DISCARD is ignored if RESTIME has not expired. If RESTIME has expired, then DISCARD is honored.

Immediate Migration Delay Interval

Table 2–18 *Immediate Migration Delay Interval*

Policy:	Values:	Default:	Command to Set Policy:
Specifies the immediate migration delay time in minutes.	0 to 9999	9999	MGMTCLAS IMMDELAY, MIGRVTV IMMDELAY

Usage Notes

- IMMDELAY lets you control the interval when immediate migration is done. For example, you may want some copies of a VTV migrated immediately while other copies are migrated at a later time, such as when the VTV meets the criteria for auto migration. The IMMDELAY parameter specifies the immediate migration delay time, which is the time in minutes to wait until immediate migrate selects VTVs for migration. Higher values preference VTV residency, which allows VTVs used in multi-step jobs to remain resident so that the jobs can complete before the VTVs are processed for migration. A value of zero equals migrate immediate.

A default strategy is to consider IMMDELAY as a switch:

- A value of 0 means migrate this VTV copy as quickly as possible.

- A value of 9999 indicates no immediate migrate (the default) and indicates that this VTV copy is migrated only when selected for auto migration.

Oracle recommends that you do not use any other values until you evaluate using IMMDELAY as a switch. Then you might want to specify some other values to “spread out” your migration activity.

- If the VTV requires replication, a value of 0 (immediate migrate) will be assigned. This can be overridden by MGMTCLAS or MIGRVTV IMMDELAY settings.

Maximum and Minimum Concurrent Migration Tasks

Table 2–19 Maximum and Minimum Concurrent Migration Tasks

Policy:	Values:	Default:	Command to Set Policy:
Specifies the maximum and minimum number of concurrent automatic migration, immediate migration, and migrate-to-threshold tasks for each VTSS.	MAXMIG - 1 to the number of RTDs connected to the VTSS. See "Maximum RTDs per VTSS" for more information. MINMIG - 1 to the MAXMIG setting.	MAXMIG - Half the number of VTDs connected to the VTSS. MINMIG - Default is 1.	MAXMIG - CONFIG VTSS MINMIG or SET MIGOPT MINMIG - CONFIG VTSS MINMIG or SET MIGOPT

Usage Notes

- Use these parameters to balance migration tasks with other tasks (such as recall and reclaim) for the RTDs you have defined for each VTSS.
- In some situations, VTCS may not be able to activate all the migration tasks specified by the MAXMIG parameter. For example:
 - The VSM-wide RTD configuration consists of four 9840 and four 9490 transports.
 - No Storage Classes with STK1R as the primary media have been defined.
 - There is sufficient MVC media for the 9490 transports.

In this configuration, because only the 9490 media is used, only a maximum of four migration tasks are activated using the 9490 RTDs.

- Similarly, there are circumstances when VTCS will start less than the number of migration tasks specified by the MINMIG parameter. For example:
 - The configuration consists of a single VTSS with 4 RTDs in ACS 0 and 4 RTDs in ACS 1. All RTD device types are identical.
 - MINMIG and MAXMIG are both set to 8.
 - Two Storage Classes are defined, which point, respectively, to ACS 0 and ACS 1.

In this configuration, if there are migrates queued for both Storage Classes, then VTCS will start 8 requests. If however, there are only migrations queued for one Storage Class, then VTCS will not start 8 requests because the workload can only be serviced by one Storage Class and this class can only run on 4 RTDs.

- Finally, also note that when you reset the MINMIG and/or MAXMIG values, the actual number of migration tasks may not be immediately affected because of the way that VTCS manages migration tasks.

Tip: You can use Display MIGRate to display migration status.

Prioritizing Automatic and Immediate Migration Workloads

Table 2–20 *Prioritizing Automatic and Immediate Migration Workloads*

Policy:	Values:	Default:	Command to Set Policy:
Prioritizes Automatic and Immediate Migration Workloads	0-99	0	MIGRSEL SCHPREF

Usage Notes

In general, MIGRSEL SCHPREF provides finer control over how Storage Classes are treated for migration. Using the MIGPOL parameter, a Management Class can request that multiple VTV copies are migrated to multiple Storage Class destinations. By default, all migrated copies are treated equally. With MIGRSEL SCHPREF, however, you can request that VTCS preferences one or more Storage Classes and the accompanying migrations.

For example, a Management Class requests two MVC copies, where one copy is business critical (perhaps for DR purposes) and the other is not. In this case, you can assign a higher preference to your business critical copy. This preferencing requests that more resources (RTDs/MVCs) are used to service the business critical copy.

MIGRSEL SCHPREF is the tool to preference Storage Classes and their accompanying migrations and MIGRSEL SCHPREF also provides the ability to set a scheduling preference for a narrow and specific set of conditions. For example, if you are migrating on a specific host from a specific VTSS and your immediate wait time has reached a specified threshold, you can then apply the specified preference value.

With that framework in mind, use MIGRSEL SCHPREF as follows:

- Use MIGRSEL SCHPREF to preference automatic and immediate migration per Storage Class. Higher values can produce quicker migration times but may not optimize MVC usage. Lower values may produce slower migration times but may optimize MVC usage.
- MIGRSEL SCHPREF allows you to specify the VTSS and host to which the preferencing applies.
- You can also specify MIGRSEL IMMWAIT, which specifies the age of the queued migrate request in minutes that the MIGRSEL rule applies to. Valid values are 0-999, the default is 999. The search criteria is considered true if the IMMWAIT value is greater than or equal to the search value.
- The MIGRSEL SCHPREF setting may be affected by the number of RTDs available, the SCHLIMIT setting, and the GLOBAL MAXMIG parameter for the VTSS.

Controlling Migration Workloads

Table 2–21 *Controlling Migration Workloads*

Policy:	Values:	Default:	Command to Set Policy:
Controls migration workloads.	0-99	99 (no limit up to the VTSS MAXMIG value)	MIGRSEL SCHLIMIT

Usage Notes

MIGRSEL SCHLIMIT, like MIGRSEL SCHPREF, provides finer control over how Storage Classes are treated for migration but in the opposite direction. That is, MIGRSEL SCHLIMIT, applies limits on migration resources (RTDs/MVCs) per Storage Class, and you can do so on a VTSS and/or host basis. Use MIGRSEL SCHLIMIT as follows:

- Use MIGRSEL SCHLIMIT to deprefer migration per Storage Class. Lower values deprefer migration, and you can specify automatic, immediate, demand, and reclaim migrates. Lower values can do the following:
 - Optimize MVC usage. You can use this strategy where you want to limit the number of MVCs used. Example: a workload migrated to a Storage Class that is subsequently used for offsite vaulting.
 - Preference migration to other Storage Classes.
 - Restricts migration concurrency to better use MVCs.
 - Limits migration to keep RTDs available for auto recalls.
 - Reduce MVC swapping when workloads change.
- MIGRSEL SCHLIMIT allows you to specify the VTSS and host to which the preferencing applies.
- For auto and immediate migration processing, MIGRSEL SCHLIMIT depreferences migration for the VTSS to Storage Class relationship. This comparison is not global and only effects requests driven by the individual VTCS host.
- For demand migration requests, MIGRSEL SCHLIMIT it will cause the request to be held if the scheduling of it would cause the number of globally active migration requests on the VTSS that satisfy the same FUNCTION and STORCLAS selection criteria to be exceeded. The migration requests will be released and a MVC picked once the constraint subsides.

Conditionally Change Migration Preferencing/Depreferencing

Table 2–22 *Conditionally Change Migration Preferencing/Depreferencing*

Policy:	Values:	Default:	Command to Set Policy:
Conditionally changes migration preferencing/depreferencing.	0-999 minutes	999	MIGRSEL IMMWAIT

Usage Notes

In the world of VTV migration, things change. VTCS puts VTV migration candidates on migration wait queues, where they wait until resources are available. Some days, VTVs are on a wait queue longer than others, so VTVs are not migrated in a timely fashion. You can use MIGRSEL IMMWAIT to conditionally manage variations in wait queue times. MIGRSEL IMMWAIT essentially sets a wait time limit, and if the wait time is less than or equal to this limit, VTCS starts the condition specified on the MIGRSEL statement to which the IMMWAIT value applies.

For example, if you have a "local" Storage Class that you want to limit/deprefer the immediate migration resources:

```
MIGRSEL STOR (LOCAL) FUNCTION (IMMED) SCHLIMIT (1)
```

You know from consulting the VTV dashboard that the normal immediate queue wait time for this Storage Class is 15 minutes. Any more than 15 minutes and your immediate migrations are no longer immediate. So to reduce some of the limiting, do the following:

```
MIGRSEL STOR (LOCAL) FUNCTION (IMMED) IMMWAIT (15) SCHLIMIT (2)
```

Again, more granularity of control: set a basic preferencing/depreferencing, and automatically adjust it based on changes in queue wait times.

Dependent Immediate Migration

Table 2–23 *Dependent Immediate Migration*

Policy:	Values:	Default:	Command to Set Policy:
Specifies dependent immediate migration policy.	INDEPEND, INITIAL, SUBSEQNT(timeout)	INDEPEND	MIGRVTV

Usage Notes

These MIGRVTV parameters Specifies dependent immediate migration policy as follows:

INDEPEND

Affected VTVs are independent of any other VTV migrate instance. This is the default and applies to any migrates that do not match a MIGRVTV rule or any rule without an INITIAL or SUBSEQNT parameter.

INITIAL

Affected VTV migrates are scheduled and performed before any migrate instances that match to a rule with the SUBSEQNT parameter.

SUBSEQNT(*timeout*)

Affected VTV migrates are scheduled and performed after any migrate instances that match to a rule with the INITIAL parameter. *timeout* specifies the time in minutes after which the rule expires. *timeout* is in addition to any IMMDELAY applied to the INITIAL migrate instances.

You can use Dependent Immediate Migration to control VLE to VLE copy. For more information, refer to *Configuring the Host Software for VLE*.

Putting It All Together: VTV Migration Control Example for an Enterprise

With VTV migration control, you have gained finer control over:

- VTSS space management.
- Speedy migration of critical data (for example, to MVCs destined for a remote ACS).
- MVC usage (for example, to ensure maximum space usage for MVCs destined for offsite vaulting).

The good part is that you can do all of the above concurrently. For example, if you want two VTV copies, one local and one remote, you can preference migration for the "remote" Storage Class **and** specify that VTVs are deleted from the VTSS ASAP after migration completes. Similarly, you can depreference migration for the "local" Storage Class and specify relatively long VTSS residency so that VTV availability is enhanced.

To demonstrate Enhanced Migration Control for an enterprise that has four typical classes of data, and name their corresponding Storage Classes accordingly:

- The REMOTE Storage Class is for MVCs going to a remote ACS. They are the primary backup for critical data so they require high migration preferencing to ensure speedy migration as described in ["Enabling the Remote and Local ACS Migrations."](#)
- The LOCAL Storage Class is for the local VTV copy of critical data. As described above, depreference migration to these MVCs (but optimize data availability) as shown in ["Enabling the Remote and Local ACS Migrations."](#)
- The LTRVLT1 Storage Class is for MVCs destined for long term retention in an offsite vaulting. You can also do vaulting for DR, and you can do both of these (and manage "floor" volumes) with the ELS Vault Solution. For more information on this scenario, refer to *Managing HSC and VTCS*. In ["Migration Policies for a Vault Solution,"](#) you are just going to show the code to achieve the above business objectives for vaulting volumes.
- The OTHER Storage Class is for temporary data sets and other things you need to safeguard, but as cheaply as possible. You want this data out of the VTSS buffer as quickly as possible, which you do using DISCARD and DELSCR as described in ["Enabling the "Other Data" Migrations."](#)

Enabling the Remote and Local ACS Migrations

In this scenario, the overall business goals are optimizing VTSS buffer use and ensuring speedy migration of critical data, while maintaining data availability. There are two MVC copies of each VTV, which are represented by these two Storage Classes:

- The REMOTE Storage Class is for MVCs going to a remote ACS. They are the primary backup for critical data so they require high migration preferencing to ensure speedy migration. The examples that follow use T10K media for the REMOTE Storage Class, so you want to optimize media use.
- The LOCAL Storage Class is for the local VTV copy of critical data. The goal here is to depreference migration to these MVCs but to optimize data availability, which includes specifying T98490D media.

And, of course, use VTSS space as efficiently as possible.

To (efficiently) enable both remote and local ACS migrations:

1. Create the LOCAL and REMOTE Storage Classes with the appropriate media and ACS values:

```
STOR NAME(LOCAL) ACS(00)MEDIA(STK1R)
STOR NAME(REMOTE) ACS(01) MEDIA(T1B000E1)
```

2. Create MIGRSEL statements corresponding to the LOCAL and REMOTE Storage Classes:

```
MIGRSEL STOR(LOCAL) FUNCTION(IMMED) SCHLIMIT(1)
MIGRSEL STOR(REMOTE) FUNCTION(IMMED) SCHPREF(9)
```

In this example:

- The LOCAL Storage Class reduces immediate migration tasks severely, because auto migration will get the data migrated soon enough.
- The REMOTE Storage Class gives immediate migration tasks the highest priority possible, because you want these VTVs on MVC space ASAP.

This step sets the priorities for the immediate migrate speed. Data availability/VTSS buffer management is addressed in the next step.

3. Create MIGRVTV statements corresponding to the LOCAL and REMOTE Storage Classes:

```
MIGRVTV STOR (LOCAL) IMMDELAY (9999)
MIGRVTV STOR (REMOTE) IMMDELAY (0)
```

These statements used the "switch" philosophy for IMMDELAY:

- REMOTE is set to 0, so migrate VTV copies to it ASAP.
- LOCAL is set to 9999, so migrate VTV copies to it only when you select the VTV for auto-migration.

As you can see, you have gained more and more granularity, finer and finer control, which continues in the next step.

4. Create a MGMTCLAS statement that points to the LOCAL and REMOTE Storage Classes:

```
MGMT NAME (COMPLY) MIGPOL (LOCAL, REMOTE) DISCARD (0) DELSCR (YES)
```

The above example provides this additional fine tuning for both Storage Classes:

- After the VTVs are successfully immediately migrated, they are not likely to be re-referenced again, so delete them from the VTSS to optimize VTSS buffer space.
- If the VTV is scratched, and you run LCM (best practices) scratch synchronization, delete the VTV from the VTSS, which is more good VTSS buffer management.

5. Create and activate an SMC POLICY statement to point to the COMPLY Management Class and specify virtual media:

```
POLICY NAME (CPLY) MEDIA (VIRTUAL) MGMT (COMPLY)
```

6. Create a TAPEREQ statement that points to the CPLY Policy:

```
TAPEREQ DSN (LOVEDONES.**) POLICY (CPLY)
```

7. Activate the new TAPEREQ set by issuing the SMC TREQDEF command.

Migration Policies for a Vault Solution

As stated in ["Putting It All Together: VTV Migration Control Example for an Enterprise,"](#) you can do vaulting for a variety of reasons and using a range of different methods. For more information on the ELS Vault Solution, refer to *Managing HSC and VTCS*. The following shows the code for the migration policies that control the data management, using a Long Term Vault (LTR) vault as an example.

To implement migration policies for vaulting:

1. Create the LTRVLT Storage Class. For example:

```
STOR NAME (LTRVLT1) ACS(00) MEDIA (STK2PB)
```

2. Create a MIGRSEL statement corresponding to the DRLOC and DRVLT Storage Classes:

```
MIGRSEL STOR (LTRVLT1) FUNCTION (IMMED) SCHPREF (9)
```

So for the LTRVLT1 Storage Class, immediate migration tasks are the highest priority possible, because you want this data protected on an MVC as soon as possible.

3. Create a MIGRVTV statement corresponding to the LTRVLT1 Storage Class:

```
MIGRVTV STOR (LTRVLT1) IMMDELAY (0)
```

LTRVLT1 IMMDELAY is set to 0, so migrate VTV copies to it as soon as possible.

4. Create a Management Class that points to the Storage Class in Step 1. For example:

```
MGMT NAME (LTRV) DISCARD(0) DELSCR (YES) MIGPOL (LTRVLT1)
```

This Management Class says:

- After the VTVs are successfully immediately migrated, they are not likely to be referenced again, so delete them from the VTSS to optimize VTSS buffer space.
 - If the VTV is scratched, and you run LCM (best practice) scratch synchronization, delete the VTV from the VTSS. This is more good VTSS buffer management.
5. Use the SMC POLICY command to point to the Management Class in Step 4. For example:

```
POLICY NAME (VLTRV1) MGMT (LTRV1)
```

6. Create a TAPEREQ statement to route the LTR data to VSM and assign the corresponding Policy/Management Class to the data. For example:

```
TAPEREQ DSN (COMPLY*) POLICY (VLTRV1)
```

7. Activate the new TAPEREQ set by issuing the SMC TREQDEF command.

Enabling the "Other Data" Migrations

The OTHER Storage Class is for temporary data sets and other things needed to safeguard, but as cheaply as possible. You want this data out of the VTSS buffer as quickly as possible, which you do using DISCARD and DELSCR.

To enable the "other data" migrations:

1. Create the OTHER Storage Class:

```
STOR NAME (OTHER) ACS (00) MEDIA (STK1R)
```

So OTHER goes to the local ACS only, 9840D media in case you need to it back up quickly. For "Other Data" Automigration runs as is, and do not worry about using MIGRSEL/MIGRVTV to fine tune things.

2. Create a MGMTCLAS statement that points to the OTHER Storage Class:

```
MGMT NAME(TEMPS) MIGPOL(OTHER) DISCARD(0) DELSCR(YES)
```

This statement specifies the following:

- After the VTVs are successfully immediately migrated, delete them from the VTSS to optimize VTSS buffer space.
 - If the VTV is scratched, and you run LCM (best practices) scratch synchronization, delete the VTV from the VTSS. This is more good VTSS buffer management.
3. Create and activate an SMC POLICY statement to point to the TEMPS Management Class and specify virtual media:


```
POLICY NAME(TDATA) MEDIA(VIRTUAL) MGMT(TEMPS)
```
 4. Create a TAPEREQ statement that points to the CPLY Policy:


```
TAPEREQ DSN(TEMP*.**) POLICY(TDATA)
```
 5. Activate the new TAPEREQ set by issuing the SMC TREQDEF command.

VTV Replication Policies

VTSS CLINKs are capable of replication, which is copying VTVs from one VTSS to another. The following sections describe VTSS replication policies.

- ["VTCS Replication - Always or Changed VTV Only."](#) Here, the concept is simple but read the section so that you understand how "changed VTVs" are defined.
- ["VTCS Replication - Synchronous or Asynchronous."](#) You have another choice with replication (assuming that you have satisfied all the requirements), which is synchronous or asynchronous processing.
- ["VTCS Enhanced Replication."](#) This method of replication starts synchronous replication on the first Read-to-Write change and copies data to the target VTSS as it is written to the sending VTSS.

VTCS Replication - Always or Changed VTV Only

Table 2–24 *VTV Replication - Always or Changed VTV Only*

Policy:	Values:	Default:	Command to Set Policy:
Specifies whether a VTV is always replicated or only when changed.	ALWAYS, CHANGED	ALWAYS	CONFIG GLOBAL REPLIcat

Usage Notes

- ALWAYS means the replicate request is added to the VTCS replication queue every time the VTV is dismounted, regardless of whether the VTV was changed while it was mounted (the default).
- CHANGED means the replicate request is added to the VTCS replication queue if the VTV:

- Was changed while it was mounted **or**
- Was only read while mounted but less than the expected number of MVC copies of the VTV exist.

Note: Regardless of the CONFIG GLOBAL REPlicat setting, replication also requires that:

- The VTV must be dismounted in a VTSS that supports replication **and** there cannot be an identical copy of the VTV in the other VTSS in the Cluster.
 - In addition to the CONFIG GLOBAL REPlicat value, you **must** specify REPlicat(YES) on a VTV's Management Class for replication to occur.
-

VTCS Replication - Synchronous or Asynchronous

Table 2–25 *VTV Replication - Synchronous or Asynchronous*

Policy:	Values:	Default:	Command to Set Policy:
Specifies whether synchronous replication is enabled.	YES, NO	NO	CONFIG GLOBAL SYNCHREP

Usage Notes

YES means that synchronous replication is enabled, but make sure you have all the requirements lined up as described in *Creating ELS Disaster Recovery Solutions*.

Note:

- Regardless of the CONFIG GLOBAL SYNCHREP setting, replication also requires that you **must** specify REPLICat(YES) on a VTV's Management Class for replication to occur.
- Synchronous replication **does incur** a performance penalty. That is, when synchronous replication is specified, the replication of the new or modified VTV to the other VTSS is made prior to the completion of the rewind/unload of the original VTV. When the virtual volume completes close processing, VTSS has completed performing the replication of the volume. Synchronous replication ensures that, under normal conditions, two identical copies of a virtual volume exist, as soon as the rewind/unload command has completed. You should consider using synchronous replication if it is mandatory for your operation to have the highest level of protection for your tape data. With synchronous replication in an unconstrained environment, the time required to replicate a virtual volume **can significantly delay** the completion of the rewind/unload command. This time will increase in a constrained environment and for larger VTV sizes.

If a synchronous replication cannot be completed immediately (likely caused by the failure of the connections between the VTSSs), the synchronous replications will operate as asynchronous replications until the failure has been corrected. If a synchronous replication is changed to an asynchronous replication, the system is notified with a warning message.

VTCS Enhanced Replication

Table 2–26 *VTV Enhanced Replication*

Policy:	Values:	Default:	Command to Set Policy:
Specifies whether enhanced replication links are enabled.	ONline, OFFline	ONline	No CONFIG changes required. Specify the Vary LINKSto command to vary all replication links to a desired state.

Usage Notes

- Virtual mount processing adds enhanced replication capabilities to all volume level synchronous requests if supported by the source VSM6/VSM7 and the destination VSM6/VSM7.

This VSM capability starts synchronous replications at the first modification of the VTV and completes the replications at Rewind Unload.

- There are no CONFIG changes required to implement support for enhanced replication. Once the supporting PTF is applied, VTCS CONFIG automatically updates the CDS with replication link support based on current CLINK definitions.

Enhanced replication requires an H level (or later) CDS.

- If CONFIG is not run, VTCS will update the CDS at startup.

Note: If enhanced replication is not supported on the VTSS, synchronous replication is done using the current method of replication starting at Rewind Unload (RUN). Virtual mount processing now selects VTSS targets for both replication link and CLINK based synchronous replications using new internal replication link structures.

- For compatibility reasons, a replication link and a CLINK(s) can be configured to the same target VTSS. The new enhanced replication requests will only be handled by the replication link and the old replication requests handled by the CLINK(s).
- Support for enhanced replication, within a local cluster or to a remote tapeplex, requires that all connectivity must support enhanced replication. Otherwise, enhanced replication is suppressed for the entire cluster or remote tapeplex.
- VTCS dismount processing incorporates new ECAM requests for querying enhanced replication attempts.

VTD Mount Failures

If the VTV mount fails, the replication link may go to an offline-fail state in the VSM6/VSM7, which is normally indicated by an ECAM 9/128 error on the mount. After this error, VTCS determines the replication link state by querying the replication link after the failure and updating its CDS status accordingly. Mount failures result in the failed replicate being done asynchronously. Once the replication link goes to offline-fail, the only way to recover the replication link is to vary it back online. Once VTCS places a replication link in offline-fail state, all replications are done asynchronously to that target until the replication link is varied back online.

VTD Dismount/VTV Audit Failures

When an enhanced replication failure occurs, the secondary VTSS automatically detects the failure, deletes the VTV, and VTCS attempts the replication asynchronously. VTCS performs normal recovery by querying the secondary VTSS and deleting the VTV if it does exist.

A replication failure does not imply a replication link state change. It is possible for a replication link connection to fail, but not all connections fail, thus the replication link stays online. VTCS does not query the replication link state as a result of a replication failure.

IMPORTING status is cleared for VTVs when the synchronous enhanced replication is complete.

Vary CLink Processing

Vary CLink processing ensures that established replication links are online before CLINKs can be brought online. Any attempt to vary a CLINK online, with an associated replication link being offline, generates an exception message, as shown in the following example.

A replication link target, DVTSS18, has been established from VTSS DVTSS16 based on existing CLINK connectivity from DVTSS16 to DVTSS18. The replication link, and the associated CLINKs, have been taken offline.

An attempt to vary an underlying CLINK online with the associated replication link offline generates a message, shown in the following example, and aborts the vary attempt.

```
Vary CLINK DVTSS16/0A - LINK DVTSS18 must be online to perform this action
```

MVC Policies

The following sections describe these MVC policies. It is usually a good idea to supplement automatic reclamation with demand reclamation as described in *Managing HSC and VTCS*.

- ["Standard Reclaim Threshold - Determines MVC Eligibility for Reclamation \(Requires Recall/Remigration\)."](#) It makes sense that you want to maximize results when you do a standard reclaim, because it costs you resources that you could be otherwise using for migrates and recalls. This is the threshold value you can use to specify how fragmented an MVC must be to make it a reclaim candidate.
- ["Dynamic Reclaim Threshold - No MVC Mount or VTV Recall/Remigration Occurs."](#) For T10000B and above media, dynamic reclaim, which requires no recalls or migrates, optimizes reclamation for capacity-centric media.

Note that the reclaim thresholds **only** dictate when an MVC becomes a reclaim candidate. They do not **start** automatic reclamation; see the next two bullets.

- ["Dynamic Reclaim - Specifying VTV Data Movement."](#) The RECLaim MOVEDATA parameter can be used to limit VTV data movement and to free up low-use partitions on an Automatically Linked Partition (ALP) MVC.

Note: Reclamation turns fragmented MVC space (space that contains non-current VTVs) into usable space (writable MVC space). MVC reports and Display MVCpool show the percentages of MVC space that is fragmented, used (space that contains current VTVs), available, and usable. Note that usable space may be zero even if there is still space physically available. For example, if the maximum VTVs per MVC is exceeded, then usable space is reported as 0%. You set maximum VTVs per MVC. Similarly, if a data check error has been reported against an MVC, VTCS will not use this MVC for output and usable space is reported as 0%.

- ["VLE Reclaim Threshold - Determines VMVC Eligibility for Reclamation."](#) There is a reclaim policy for VMVCs.
- ["Free MVCs Threshold - Starts Automatic Space Reclamation"](#) and ["Eligible/Total MVCs Threshold - Starts Automatic Space Reclamation."](#)

There are two different triggers that start automatic reclamation. More fine tuning, more granularity. These triggers supplement rather than compete with each other. If either you run low on Free MVCs or your ratio of eligible to total MVCs gets too low, automatic reclamation starts.

- ["Maximum MVCs Processed Per Reclaim."](#) VTCS reclaims MVCs one at a time. You can, however, specify the maximum number you want to reclaim in one run, either automatic or demand. You can specify that when reclamation starts, run it to reclaim this number of MVCs, because that is the value that allows you to have sufficient free MVCs.

- ["Maximum MVCs Concurrently Processed for Reclamation and Drain."](#) A reasonable question is, what is the difference between this and maximum MVCs processed per reclaim, especially since you can only reclaim one MVC at a time? It turns out that reclaim and drain are separate but related processes that can be processing multiple MVCs, and you get to say how many.
- ["Hosts Disabled from Reclamation."](#) If you do not want a host to initiate reclamation, you can explicitly turn it off. The default is not to disable any hosts, so if this is not an issue, do not worry about it.
- ["MVC Retain Interval."](#) The idea is pretty simple: if you think you are going to write more data to an MVC once it is mounted and initially written to, you can specify a value that keeps it available. A similar control is the MVC mount timeout interval, see ["MVC Mount Timeout Interval."](#)
- ["Non-Library Resident MVC Policies."](#) The non-library resident MVC policies specify how VTCS processes MVCs that are not in the ACS.
- ["MVC Initialization on First Mount."](#) Initializing MVCs before they are used is no longer required, thanks to this global and MVC Pool parameter. You can now specify that the MVC is initialized at first mount.

Standard Reclaim Threshold - Determines MVC Eligibility for Reclamation (Requires Recal/Remigration)

Table 2–27 Standard Reclaim Threshold

Policy:	Values:	Default:	Command to Set Policy:
Specifies the standard reclaim threshold (as a percentage) that determines when an MVC is eligible for demand or automatic reclamation.	4 - 98	40	CONFIG RECLAIM THRESHLD RECLAIM THRESHLD MVCPOOL THRESHP

Usage Notes

- Inevitably, MVCs become separated with gaps where non-current VTVs reside, which is sub-optimal space use. Customers can use the CONFIG utility to set MVC space reclamation parameters, and they also can use the RECLAIM command or utility to adjust these parameters. When MVCs reach specified thresholds, VTCS automatically starts standard space reclamation, reclaiming space one MVC at a time. VSM reclaims space by copying only valid VTVs from the selected MVC to the VTSS, then copying these VTVs back to another MVC with sufficient space. VSM copies only those VTVs placed after the first open space on the MVC, which saves I/O cycles. The space reclamation thus reduces MVC fragmentation and allows the migration to succeed.

If fragmented space on an MVC exceeds the value specified on THRESHLD, VTCS makes the MVC eligible for reclamation. Regardless of the percentage of fragmented space on an MVC versus this value, however, VTCS also considers where fragmented space occurs. For example, if the first fragmented space is near the end of the MVC, VTCS may process the MVC before an MVC with more total fragmented space.

- You can use Display MVCpool to display the MVCs eligible for reclamation in your MVC pool, as well as information about MVC status and space.

- To optimize recall processing in mixed-media systems, ensure that your MVC pool has at least one media type compatible with each RTD type.
- The MEDIA parameter of the STORCLAS statement (which defines Storage Classes) specifies a preference list of MVC media types. This list supersedes the default media selection list.
- If the system's MVCs are highly fragmented, a demand standard MVC space reclamation can be scheduled (using RECLAIM) as an off-hours batch job.
- Dynamic reclaim does not apply to T1000D Sport cartridges. Dynamic reclaim requires an H level CDS, VSM5 microcode D02.11.16.00 or later, T10KB firmware 1.41a.209 or later.
- You can specify Dynamic Reclaim policies at the global, system wide level or on a Storage Class and MVC pool level. The global policies are set by the VTCS CONFIG RECLAIM statement:
 - The INPLACE(YES/NO) parameter controls the usage of dynamic reclaim. The INPLACE setting can either be explicitly specified upon the STORCLAS statement or inherited from the RECLAIM statement in the VTCS configuration. The default is INPLACE=NO. The setting that is derived for a Storage Class does not force usage of media formatted in either mode. So it is quite acceptable to use with drives or media that do not support partitioning. Also, any media written in a different mode will continue to be a migration target. This will avoid the need to actively manage the transition when dynamic reclaim is activated. Its effect is to cause new media allocated to the Storage Class to be formatted and written in either standard (INPLACE=NO) or partitioned (INPLACE=YES) mode. The switch can only occur when media is empty.

If a hard switch-over is required between standard and dynamic reclaim, then a lengthy process is required. Parameters setting (for example, MEDIA on the STORCLAS statement) will need to be used. MVCs will need to be either taken out of the MVC pool or set into a read-only state by using MVCMAINT. Finally, MVCs formatted in the wrong mode would need to be drained.
 - The INPTHRSH parameter specifies the fragmentation level that makes partitioned MVCs eligible for Dynamic reclaim. When a partitioned MVC's fragmentation falls between the INPTHRSH value and the THRESH value, it first becomes eligible for Dynamic Reclaim. If no space is freed on the initial reclaim attempt, the reclaim will be retried when the fragmentation level has risen significantly. Should the fragmentation level reach the standard THRESHOLD be reached, then a fall-back method of moving VTVs will be performed. This will still be less than for standard format media.
- The VTCS CONFIG RECLAIM statements provide the global defaults for the system. These can be overridden at the Storage Class and MVCPool level:
 - The STORCLAS INPLACE parameter specifies whether or not Dynamic Reclaim is enabled for this Storage Class. Note that initially, you may want to specify INPLACE(NO) in the VTCS CONFIG RECLAIM statement and INPLACE(YES) for selected Storage Classes so you can evaluate the effects of Dynamic Reclaim. Subsequently, if you decide you want Dynamic Reclaim to apply to all eligible media, you can specify INPLACE(YES) on the VTCS CONFIG RECLAIM statement.
 - The INPTHRSH parameter on the POOLPARM statement can be used to set a Dynamic Reclaim threshold specifically for this set of MVC media. This can be

used if you wish to treat some MVCs differently from the VTCS CONFIG RECLAIM defaults.

- Dynamic Reclaim can be disabled either globally at a system wide level or at the Storage Class level:
 - To disable Dynamic reclaim globally specify INPLACE(NO) in the VTCS CONFIG RECLAIM statement and ensure that any Storage Class definition does not override this to YES.
 - To disable Dynamic for one or more Storage Classes change INPLACE to NO in the Storage Class policy statements.

Note: Any MVCs that are currently in use for migration will continue in partitioned format. If you wish to stop active MVCs from further migrates then use the MVCMAINT utility to mark them as read-only. Once the partitioned volumes are empty they will be automatically re-used in non-partitioned mode. An MVC will become empty either when all VTVs have expired or the MVC has been DRAINEd.

Dynamic Reclaim Threshold - No MVC Mount or VTV Recall/Remigration Occurs

Table 2–28 Dynamic Reclaim Threshold

Policy:	Values:	Default:	Command to Set Policy:
Specifies if dynamic reclaim is enabled.	YES, NO	No default	CONFIG RECLAIM INPLACE STORCLAS INPLACE
Specifies the dynamic reclaim threshold (as a percentage) that determines when an MVC is eligible for demand or automatic reclamation.	Less than CONFIG RECLAIM THRESHLD and between 3 and 97	Half the CONFIG RECLAIM THRESHLD value (rounded up).	CONFIG RECLAIM INPTHRSH MVCPool INPTHRSH POOLPARAM TYPE(MVC) INPTHRSH

What is Dynamic Reclaim?

As shown in [Figure 2–2](#), with a standard MVC whose fragmentation exceeds the standard reclaim threshold, standard reclaim must recall/remigrate VTVs (V3, V4, and V5, in this example) in order to reclaim MVC space.

Figure 2–2 Standard Reclaim Example



[Figure 2–3](#) shows a blank partitioned MVC with all partitions free.

Figure 2–3 Blank Partitioned MVC



Figure 2–4 shows a partitioned MVC with all partitions full. The VTVs are in light blue, on top, and the partitions are in dark blue, on the bottom.

Figure 2–4 Full Partitioned MVC



Figure 2–5 shows a fragmented partitioned MVC, where VTVs 6 and 14 have expired, which frees partitions 6, 7, 15, and 16. These partitions are "dynamically reclaimed," which means they are now available for migration of new VTVs. For dynamic reclaim, VTCS just updates the CDS to mark the freed partitions as writeable; the MVC is not mounted and no data movement occurs.

Figure 2–5 Dynamically Reclaimed Partitioned MVC



Figure 2–6 shows the dynamically reclaimed MVC shown in Figure 2–5. VTVs 21 and 22 were subsequently migrated to partitions 6 and 7 and 15 and 16 respectively. Note that VTV 21 spans partitions 6 and 7 and VTV 22 spans partitions 15 and 16 because each partition on the MVC is automatically linked to the next partition.

Figure 2–6 Dynamically Reclaimed Partitioned MVC - New Migrations



Usage Notes

- Dynamic reclaim only applies to T10000B and above media (full and full encrypted only) that have been formatted by T10000 drives capable of creating partitions.
- For 32 GB VSM 6 VTVs, eight VTVs can span a partition.

Dynamic Reclaim - Specifying VTV Data Movement

The RECLaim command performs demand MVC space reclamation. Use the MOVEDATA parameter to limit VTV data movement and to free up low-use partitions on an Automatically Linked Partition (ALP) MVC.

Table 2–29 Dynamic Reclaim - Data Movement

Command:	Values:	Default:	Command to Set Policy:
Specifies the extent of VTV data movement that results from an MVC being reclaimed.	Standard, Minimum, None	Standard	RECLaim MOVEDATA

Usage Notes

- To reclaim only your eligible ALP MVCs INPLACE and ensure no data movement occurs for MVCs over the THRESH value, include the parameter MOVEDATA(None) on associated dynamic reclaim commands.

Figure 2–7 shows an MVC report with three ALP MVCs.

Figure 2–7 Three ALP MVCs and MOVEDATA(None)

```

TIME 10:22:40
MVC Number
Volser VTVs %Used %Avail %Frag Size Mount Status Last Mounted Location Owner/Comment/
(GB) Count J B L D R U T M Date Time VTSS / ACS ID Consolidate Time
DMV600 10 4.39 90.76 4.85 930.0 257 R - - - U - 2015Sep16 10:13:29 DVTSS16_00 S0
DMV601 24 5.27 89.22 5.51 930.0 257 R - - - U - 2015Sep16 10:13:29 DVTSS16_00 S0
DMV602 83 9.11 82.73 8.16 930.0 257 R - - - U - 2015Sep16 10:13:29 DVTSS16_00 S0
    
```

When MOVEDATA(NONE) is specified:

```
RECLAIM STORCLAS(S0) INPTHRSH(3) THRESHLD(6) MOVEDATA(NONE)
```

MVC DMV600 contains VTVs with changed MGMTclas definitions specifying different STORclas(es). It is above the INPTHRSH value and below the THRESHLD value. It has empty partitions. It is eligible for INPLACE reclaim.

MVC DMV601 is above the INPTHRSH value and below the THRESHLD value. It has empty partitions. It is eligible for INPLACE reclaim. There are no VTVs with MGMTclas definitions requiring reconciliation, movement, or other actions.

MVC DMV602 is above the THRESHLD value.

After specifying MOVEDATA(NONE), Figure 2–8 shows the MVC report after reclaim.

Figure 2–8 MVC Report After MOVEDATA(None) Specified

```

TIME 10:41:22
MVC Number
Volser VTVs %Used %Avail %Frag Size Mount Status Last Mounted Location Owner/Comment/
(GB) Count J B L D R U T M Date Time VTSS / ACS ID Consolidate Time
DMV600 10 4.40 94.42 1.18 930.0 257 R - - - U - 2015Sep16 10:13:29 DVTSS16_00 S0
DMV601 24 5.28 89.84 4.88 930.0 257 R - - - U - 2015Sep16 10:13:29 DVTSS16_00 S0
DMV602 83 9.13 81.72 9.15 930.0 257 R - - - U - 2015Sep16 10:13:29 DVTSS16_00 S0
    
```

MVC DMV600 contains the same number of VTVs. No data movement has occurred for the VTVs with changed MGMTclas or STORclas definitions. The fragmentation percentage has been reduced because a number of partitions have been freed up by INPLACE.

MVC DMV601 contains the same number of VTVs. No data movement has occurred, but some space has been reclaimed from empty partitions freed by INPLACE.

MVC DMV602 contains the same number of VTVs. No data movement has occurred even though the fragmentation was over the THRESHLD value. Some space has been reclaimed from empty partitions with the EOT position reset. However, in this particular test case, due to the layout of data across partitions on the MVC, the %Used and %Frag calculated values have actually increased slightly.

- To reclaim your eligible ALP MVCs INPLACE and free up partitions with free space exceeding the THRESH value, include the parameter MOVEDATA(Minimum) on associated dynamic reclaim commands.

Figure 2–9 shows an MVC Report with two ALP MVCs.

Figure 2–9 Two ALP MVCs and MOVEDATA(Minimum)

```

TIME 04:28:46                                VTCS MVC Summary Report                                DATE 2015-10-31
MVC      Number                               Size Mount-----Status-----Last Mounted-----> Location Owner/Comment/
Volser  VTVs %Used %Avail %Frag (GB) Count I B L D R U T M Date Time VTSS / ACS ID Consolidate Time
DMV600  212  14.16 65.05 20.79 230.0 258 R - - - - U - 2015Oct31 04:27:28 DVTSS16 00 S0
DMV601  238  31.29  9.96 58.75 230.0 258 R - - - - F - 2015Oct31 04:22:10 DVTSS16 00 S0
...
212 VTVs found for MVC:DMV600      Last migrate:2015Oct31 04:27:28
117 unused / 75 used partitions found for MVC:DMV600
Unused partitions:
-----
04B 04C 04D 04E 04F 050 051 052 053 054 055 056 057 058 059
05A 05B 05C 05D 05E 05F 060 061 062 063 064 065 066 067 068 069 06A 06B 06C 06D 06E 06F 070 071 072 073 074 075 076 077
078 079 07A 07B 07C 07D 07E 07F 080 081 082 083 084 085 086 087 088 089 08A 08B 08C 08D 08E 08F 090 091 092 093 094 095
096 097 098 099 09A 09B 09C 09D 09E 09F 0A0 0A1 0A2 0A3 0A4 0A5 0A6 0A7 0A8 0A9 0AA 0AB 0AC 0AD 0AE 0AF 0B0 0B1 0B2 0B3
0B4 0B5 0B6 0B7 0B8 0B9 0BA 0BB 0BC 0BD 0BE 0BF
...
238 VTVs found for MVC:DMV601      Last migrate:2015Oct31 04:22:10
6 unused / 186 used partitions found for MVC:DMV601
Unused partitions:
-----
0BA 0BB 0BC 0BD 0BE 0BF

```

When MOVEDATA(Minimum) is specified:

```
RECLAIM STORCLAS(S0) INPTHRSH(5) THRESHLD(90) MOVEDATA(MINIMUM)
```

MVC DMV600 contains VTVs of varying sizes. It is above the INPTHRSH value and below the THRESHLD value. It does not contain any empty partitions, but it does have some low-use partitions. It is eligible for INPLACE reclaim and some partitions can be expected to be freed by defrag.

MVC DMV601 contains VTVs of varying sizes. It is above the INPTHRSH value and below the THRESHLD value. It does not have any empty partitions but it does have some low-use partitions. It is eligible for INPLACE reclaim and some partitions can be expected to be freed by defrag.

After specifying MOVEDATA(Minimum), Figure 2–10 shows the MVC report after reclaim.

Figure 2–10 MVC Report After MOVEDATA(Minimum) Specified

```

TIME 04:57:50
MVC Number
Volser VTVs %Used %Avail %Frag Size Mount Status Last Mounted Location Owner/Comment/
DMV600 203 14.06 64.94 21.00 259 R - - - U - 2015oct31 04:40:30 DWTSS16_00 S0
DMV601 113 25.89 69.01 5.10 259 R - - - U - 2015oct31 04:40:31 DWTSS16_00 S0
DMV602 125 5.48 94.52 0.00 257 R - - - U - 2015oct31 04:56:17 DWTSS16_00 S0
DMV603 9 0.13 99.87 0.00 257 R - - - U - 2015oct31 04:43:28 DWTSS16_00 S0
...
203 VTVs found for MVC:DMV600 Last migrate:2015oct31 04:27:28
123 unused / 69 used partitions found for MVC:DMV600
Unused partitions
000 001 002 003 --- --- 006 --- --- --- --- --- --- --- --- --- --- --- --- --- --- ---
043 --- --- --- --- --- --- --- --- --- --- --- --- --- --- --- --- --- --- --- ---
05A 05B 05C 05D 05E 05F 060 061 062 063 064 065 066 067 068 069 06A 06B 06C 06D 06E 06F 070 071 072 073 074 075 076 077
078 079 07A 07B 07C 07D 07E 07F 080 081 082 083 084 085 086 087 088 089 08A 08B 08C 08D 08E 08F 090 091 092 093 094 095
096 097 098 099 09A 09B 09C 09D 09E 09F 0A0 0A1 0A2 0A3 0A4 0A5 0A6 0A7 0A8 0A9 0AA 0AB 0AC 0AD 0AE 0AF 0B0 0B1 0B2 0B3
0B4 0B5 0B6 0B7 0B8 0B9 0BA 0BB 0BC 0BD 0BE 0BF
...
113 VTVs found for MVC:DMV601 Last migrate:2015oct31 04:22:10
131 unused / 61 used partitions found for MVC:DMV601
Unused partitions
000 --- --- --- --- --- --- 007 008 009 00A 00B 00C --- --- 00F 010 011 012 013 014 015 016 017 018 019 01A 01B 01C 01D
01E 01F 020 021 022 023 024 025 026 027 028 029 02A 02B 02C 02D 02E 02F 030 031 032 --- --- --- --- --- --- --- --- ---
--- --- --- --- --- --- --- --- --- --- --- --- --- --- --- --- --- --- --- --- --- --- --- --- --- --- --- ---
065 066 067 068 069 06A 06B 06C 06D 06E 06F 070 071 072 073 074 075 076 077
078 079 07A 07B 07C 07D 07E 07F 080 081 082 083 084 085 086 087 088 089 08A 08B 08C 08D 08E 08F 090 091 092 093 094 095
096 097 098 099 09A 09B 09C 09D 09E 09F 0A0 0A1 0A2 0A3 0A4 0A5 0A6 0A7 0A8 0A9 0AA 0AB 0AC 0AD 0AE 0AF 0B0 0B1 0B2 0B3
0B4 0B5 --- 0B7 --- --- 0BA 0BB 0BC 0BD 0BE 0BF

```

MVC DMV600 has had 9 VTVs moved off it by defrag. The used partition count has gone from 75 to 69, that is, 6 partitions have become free. Some space has been reclaimed from these defragged partitions and the EOT position has been reset. However, in this particular test case, because of the layout of data across partitions on the MVC, the %Used, %Avail and %Frag calculated values have not changed much, even though a number of partitions were freed.

MVC DMV601 has had 125 VTVs moved off it by defrag. The used partition count has gone from 186 to 61, that is, 125 partitions have become free. A lot of space has been reclaimed from these defragged partitions and the EOT position has been reset. In this particular case, because of the layout of data across partitions on the MVC, the %Used, %Avail and %Frag calculated values have changed significantly.

The partition size for a T10KB MVC is 4727MB. The following VTVs demonstrate the defrag THRESH of 90% in the command selecting appropriately sized VTVs to move.

A THRESH of 90% for a 4727MB partition gives a VTV used space threshold size of 473MB, that is, VTVs < 473MB will be subject to defrag. Calculated VTV sizes are rounded up by 1MB.

The VTVs in Figure 2–11 existed on DMV601 before the RECLAIM command was issued:

Figure 2–11 VTVs Before Reclaim

DX1012	409.60	0005/000c1c59-0006	M5	128008	32008	empty space	to previous
DX1014	409.60	0006/000f0a69	M5	12808	64008	empty space	to previous
DX1016	409.60	0007/00113079	M5	12808	26521	empty space	to previous
DX1018	409.60	0008/0012d989	M5	12808	35337	empty space	to previous
DX1020	409.60	0009/0014ff99	M5	12808	76153	empty space	to previous
DX1022	409.60	000A/001725A9	M5	12808	16969	empty space	to previous
DX1024	409.60	000B/00194BB9	M5	12808	57785	empty space	to previous
DX4000	471.04	000C/00197dc1	M5	14728	70593	empty space	to previous
DX4003	472.06	000D/001BEE99	M5	14760	18265	empty space	to previous
DX4006	473.08	000E/001DFF71	M5	14792	65937	empty space	to previous
DX1028	409.60	000F/00223451	M5	12808	41617	empty space	to previous

After the RECLAIM has completed, the same section of the detail report now appears as shown in [Figure 2–12](#).

Figure 2–12 VTVs After Reclaim

```

DX1012      4096.00 0005/000c1c59-0006   M5           128008  32008 empty space to previous
DX1014      409.60 0006/000f0a69         M5           12808   64008 empty space to previous
DX4003      472.06 000D/001BEE99         M5           14760   18265 empty space to previous
DX4006      473.08 000E/001DFF71         M5           14792   65937 empty space to previous
DX1099      4096.00 0033/006e9a89         M5           128008  82985 empty space to previous

```

VTV DX1014, on partition 0006, is not eligible for defrag because the previous VTV, DX1012, also spans into partition 0006. The total used space of the two VTVs, which needs to be recalled to free the partition, is greater than the threshold size of 473MB for the partition.

VTVs DX1016-DX1024 were all moved off by defrag because their size of 409.6MB, rounded up to 410MB, is less than the threshold used space of 473MB for the partition.

VTV DX4000 was moved off by defrag because its size is 471.04MB, rounded up to 472MB, is less than the threshold size of 473MB for the partition.

VTV DX4003 was not eligible for defrag because its size is 472.06MB, rounded up to 473MB, is not less than the threshold size of 473MB for the partition.

VTV DX4006 was not eligible for defrag because its size is 473.08MB, rounded up to 474MB, is not less than the threshold size of 473MB for the partition.

VTVs DX1028-DX1098 were all moved off by defrag because their size of 409.6MB, rounded up to 410MB, is less than the threshold used space of 473MB for the partition.

- To stop data movement occurring from VMVCs in a VLE, specify `MOVEDATA(Minimum)` or `MOVEDATA(None)` on dynamic reclaim commands targeting Storage Classes in their VLE(s).

The following example shows an MVC Report with three VLE MVCs.

Figure 2–13 Three VLE MVCs with MOVEDATA(Minimum)

```

TIME 06:21:54
MVC Number                               VTCS MVC Summary Report                               DATE 2015-09-09
Volser  VTVs  %Used  %Avail  %Frag  Size  Mount  Status  Last Mounted  Location Owner/Comment/
VMC000  10  16.38  65.60  18.02  250.0  257  I - - - - U - 2015Sep09 06:18:06 DVTSS16 VLELIB0 R0
VMC001  24  19.66  59.86  20.48  250.0  257  I - - - - U - 2015Sep09 06:18:06 DVTSS16 VLELIB0 R0
VMC002  83  33.99  35.70  30.31  250.0  257  I - - - - U - 2015Sep09 06:18:07 DVTSS16 VLELIB0 R0

```

When `MOVEDATA(Minimum)` is specified:

```
RECLAIM STORCLAS(S0) THRESHLD(15) MOVEDATA(MINIMUM)
```

VMVC VMC000 contains VTVs with changed `MGMTclas` definitions specifying different `STORclas(es)`. It is above the `THRESHLD` value. It has empty partitions.

VMVC VMC001 is above the `THRESHLD` value. It has empty partitions. There are no VTVs with `MGMTclas` definitions requiring reconciliation, movement, or other actions.

VMVC VMC002 is above the `THRESHLD` value.

After specifying MOVEDATA(Minimum), the MVC report appears as shown in Figure 2–14.

Figure 2–14 VLE MVC Report After MOVEDATA(Minimum) Specified

```

TIME 06:25:09                               VTCS MVC Summary Report                               DATE 2015-09-09
MVC      Number                               Size Mount -----Status----- Last Mounted-----> Location Owner/Comment/
volser   VTSS %Used %Avail %Frag (GB) Count I B L D R U T M Date   Time   VTSS / ACS ID Consolidate Time
VMC000   10  16.38 83.62  0.00 250.0 258 I - - - - U - 2015Sep09 06:22:32 DVTSS16 VLELIBQ R0
VMC001   24  19.66 80.34  0.00 250.0 257 I - - - - U - 2015Sep09 06:18:06 DVTSS16 VLELIBQ R0
VMC002   83  33.99 66.01  0.00 250.0 257 I - - - - U - 2015Sep09 06:18:07 DVTSS16 VLELIBQ R0
    
```

Each VMVC has had all the space reclaimed and no data movement has occurred for any of them due to the MOVEDATA(Minimum) parameter.

For VLE VMVCs the same result would have occurred if MOVEDATA(None) had been specified.

VLE Reclaim Threshold - Determines VMVC Eligibility for Reclamation

Table 2–30 VLE Reclaim Threshold

Policy:	Values:	Default:	Command to Set Policy:
Specifies the reclaim threshold (as a percentage) that determines when an VMVC is eligible for demand or automatic reclamation.	4 - 98	30	CONFIG RECLAIM VLTHRES

Usage Notes

- CONFIG RECLAIM VLTHRES applies only to VMVCs.
- VLE MVC media (VMVCs) is subject to fragmentation and must be reclaimed just like real MVCs. The VMVC reclaim process, however, uses far fewer resources than a standard reclaim. The reclaim threshold for a VMVC is specified through the CONFIG RECLAIM VLTHRES parameter. The lower that you set VLTHRES, the more frequent VTCS will run reclaim on the VMVCs and the greater the effective capacity of the VMVCs (less fragmentation).

Free MVCs Threshold - Starts Automatic Space Reclamation

Table 2–31 Free MVCs Threshold

Policy:	Values:	Default:	Command to Set Policy:
Specifies the minimum number of free MVCs. If the actual number drops below this value, automatic reclamation starts.	0 - 255	40	CONFIG GLOBAL MVCFREE

Usage Notes

- A free MVC has 100% usable space and does not contain any migrated VTVs.
- VTCS checks the MVCFREE value for each ACS. VTCS issues message SLS6616I and starts an automatic space reclamation if **both** of the following occurs:
 - Free MVCs is equal to or less than the value specified on CONFIG MVCFREE.
 - There is at least one eligible MVC as defined by the CONFIG RECLAIM THRESHLD parameter; for more information, see "[Eligible/Total MVCs Threshold - Starts Automatic Space Reclamation.](#)"
- If you set MVCFREE=0, VTCS actually uses the default value (40).
- Oracle **strongly recommends** that you ensure that your MVC pool always has at least one eligible MVC for each MVC media type.

Otherwise, you may need to change the CONFIG GLOBAL MVCFREE value, add more MVCs to the pool, or both. You can use Display to display the number of free MVCs in your MVC pool.

Eligible/Total MVCs Threshold - Starts Automatic Space Reclamation

Table 2–32 Eligible/Total MVCs Threshold

Policy:	Values:	Default:	Command to Set Policy:
Specifies a percentage value that represents the ratio of reclaim candidates to total MVCs, which triggers automatic space reclamation as described in " Usage Notes. "	1 - 98	35	CONFIG RECLAIM START

Usage Notes

- CONFIG RECLAIM START specifies a percentage value, which is equal to:

$$(\text{Reclaim Candidates} / \text{Reclaim Candidates} + \text{Free MVCs}) * 100$$

Where:

Reclaim Candidates

The number of Reclaim Candidates determined by the CONFIG RECLAIM THRESHLD parameter. For more information, see "[Eligible/Total MVCs Threshold - Starts Automatic Space Reclamation.](#)"

Reclaim Candidates + Free MVCs

Equals the number of Reclaim Candidates plus the number of free MVCs. A free MVC:

- has 100% usable space and does not contain any migrated VTVs.
- is defined as described in "[Defining MVCs.](#)"
- is writeable.
- is resident in the ACS.

- For each ACS (not globally for all ACSs), VTCS issues message SLS6616I and starts an automatic space reclamation if both of the following occurs:
 - The actual value of (Reclaim Candidates/Reclaim Candidates + Free MVCs) * 100 exceeds the value specified on CONFIG RECLAIM START parameter.
 - The number of eligible MVCs exceeds the value specified on the MAXMVC parameter; for more information, see ["Maximum MVCs Processed Per Reclaim."](#)

Note: The only exception to the above two conditions occurs if an SLS6699 message indicates a critical shortage of free MVCs, in which case automatic reclamation will start anyway.

- The following are general guidelines for specifying values for the START parameter:
 - A low value (for example, 5%), starts automatic space reclamation when there are few eligible MVCs compared to free MVCs **unless** you set the MAXMVC value high compared to the number of eligible MVCs.
 - A high value (for example, 95%), starts automatic space reclamation when there are **many** eligible MVCs compared to free MVCs unless you set the MAXMVC value **very** high and your MVC pool is **very** small.
- You can use Display MVCPOOL to display eligible and free MVCs.

Maximum MVCs Processed Per Reclaim

Table 2–33 *Maximum MVCs Processed Per Reclaim*

Policy:	Values:	Default:	Command to Set Policy:
Specifies the maximum number of MVCs that will be processed in a single space reclamation run.	1 - 98	40	CONFIG RECLAIM MAXMVC RECLAIM MAXMVC

Usage Notes

- Automatic and demand space reclamation processes one MVC at a time. You can, however, use MAXMVC to control the **maximum** number of MVCs that will be processed in a single space reclamation run (automatic or demand).
- For automatic space reclamation to start using the CONFIG RECLAIM START parameter setting, the number of eligible MVCs (determined by the CONFIG RECLAIM THRESHLD parameter) must also exceed the MAXMVC value. For more information, see ["Eligible/Total MVCs Threshold - Starts Automatic Space Reclamation."](#)
- The following are general guidelines for specifying values for the MAXMVC parameter:
 - A **low** value reclaims **fewer** MVCs in a single run but may have negligible effect on migrations and recalls and may start automatic space reclamation **more** frequently; for more information, see ["Eligible/Total MVCs Threshold - Starts Automatic Space Reclamation."](#)

- A high value reclaims **more** MVCs in a single run but may have **considerable** effect on migrations and recalls and may start automatic space reclamation **less** frequently; for more information, see "[Eligible/Total MVCs Threshold - Starts Automatic Space Reclamation.](#)"
- You can use Display MVCpool to display eligible and free MVCs.

Maximum MVCs Concurrently Processed for Reclamation and Drain

Table 2–34 shows maximum MVCs concurrently processed for reclamation and drain.

Table 2–34 *Maximum MVCs Concurrently Processed for Reclamation and Drain*

Policy:	Values:	Default:	Command to Set Policy:
Specifies the maximum number of MVCs concurrently processed for reclamation and drain.	1 - 99	1	<ul style="list-style-type: none"> ■ CONFIG RECLAIM CONMVC ■ MVCDRAIN CONMVC ■ RECLAIM CONMVC

Hosts Disabled from Reclamation

Table 2–35 shows hosts disabled from reclamation.

Table 2–35 *Hosts Disabled from Reclamation*

Policy:	Values:	Default:	Command to Set Policy:
Specifies that a host cannot initiate automatic and demand reclamation.	NORECLAM, if not specified, reclamation is not disabled	Do not disable	CONFIG HOST NORECLAM

Usage Notes

- Specifying NOMIGRAT also causes NORECLAM to be set.
- Disabled hosts can still do demand MVC drains using MVCDRAIN.

MVC Retain Interval

Table 2–36 *MVC Retain Interval*

Policy:	Values:	Default:	Command to Set Policy:
Specifies how long VTCS will retain an MVC on an RTD in idle mode after a migration.	1 - 60 minutes	10	CONFIG VTSS RETAIN parameter

Usage Notes

- Retaining the MVC can reduce MVC mounts.

- When VTCS shuts down, VTCS dismounts all MVCs regardless of the MVC retain interval.

MVC Mount Timeout Interval

Table 2–37 MVC Mount Timeout Interval

Policy:	Values:	Default:	Command to Set Policy:
Specifies the interval before MVC mount timeout.	5 - 30 minutes	15	CONFIG GLOBAL MVMNTTO= <i>mm</i> parameter

Non-Library Resident MVC Policies

The non-library resident MVC policies specify how VTCS processes MVCs that are not in the ACS.

Table 2–38 Non-Library Resident MVC Policies

Policy:	Values:	Default:	Command to Set Policy:
Specifies how VTCS handles the mount of non-library resident MVCs during drain/reclaim processing.	YES (VTCS requests the mount of the non-library MVC), NO (VTCS suppresses the mount and purges the request). Note: For reclaim, only library resident MVCs can be selected for processing, never non-library ones. However, between the time a library resident MVC is selected and the time it is actually processed, it may have become non-resident by being ejected. For drain, non-library resident MVCs can be selected.	YES	CONFIG GLOBAL NLIBDRNR
Specifies whether non-library resident MVCs are selected for migration processing.	YES, NO	YES (select non-library resident MVCs for migration)	CONFIG GLOBAL NLIBMIGR
Specifies whether non-library resident MVCs are selected for reclamation processing.	YES, NO	YES (select non-library resident MVCs for reclamation)	CONFIG GLOBAL NLIBRECL

MVC Initialization on First Mount

Table 2–39 MVC Initialization on First Mount

Policy:	Values:	Default:	Command to Set Policy:
Specifies whether uninitialized MVCs are initialized when first mounted.	NO, YES	NO for CONFIG GLOBAL, none for MVCPOOL	CONFIG GLOBAL INITMVC parameter, MVCPOOL NAME INITMVC parameter

Note:

- MVCPOOL INITMVC overrides GLOBAL INITMVC. There is no default for MVCPOOL INITMVC; if not specified for a named MVC Pool the CONFIG GLOBAL value (or default) is used.
 - Initialization of MVCs in the DEFAULTPOOL is controlled by the GLOBAL INITMVC specification (or default).
 - MVC Initialization applies only to VSM4/5 and requires microcode level D02.05.00.00 or higher. If this level of microcode is not installed on all VTSSs in the configuration, MVC initialization will be limited to the VTSSs that have it installed.
-

Preparing for Configuration

In this chapter, you start configuring the ELS software. But first, there are several items you need to get done in MVS to ensure a successful configuration. So before configuring HSC/VTCS, complete the preparation tasks described in the following sections:

- ["Defining A Security System User ID for HSC and VTCS"](#)
- ["Configuring MVS Device Numbers and Esoterics"](#)
- ["Setting the MVS Missing Interrupt Handler \(MIH\) Value"](#)

Note: Refer to *Installing ELS* for information on installing the ELS software.

Defining A Security System User ID for HSC and VTCS

ELS software (HSC and VTCS) uses the MVS System Authorization Facility (SAF) to control the usage of MVCs. Before mounting an MVC and before issuing writes to an MVC, VSM issues a SAF query to verify that the HSC user has UPDATE authority for the MVC.

You must define a security system user ID to be used by the SAF queries. All VSM SAF queries are issued by HSC, so the security system user ID must be associated with the HSC started task. Refer to your security system documentation for details on how to associate a security system user ID with the HSC started task.

You must also set up tape volume security profiles to ensure that VSM is authorized to mount MVCs, to ensure that VSM is authorized to update MVCs, and to guard against accidental overwrites of MVCs; for more information, see ["Defining MVC Pool Volser Authority."](#)

Caution: Depending on the default settings of your security system, VSM may not be able to mount and to write to MVCs until you have defined a security system user ID for HSC and tape volume security profiles for the MVCs.

Configuring MVS Device Numbers and Esoterics

The following sections tell how to use the HCD facility to do the following:

- Assign MVS device numbers to VTDs, library drives, and RTDs.
- Associate these MVS device numbers and esoterics.

You determined these values in "[Determining HSC/VTCS Configuration Values.](#)" Refer to your IBM documentation for more information on the HCD facility.

Assigning MVS Device Numbers to VTDs and Library Drives

Use the HCD facility to assign MVS device numbers to these devices (VTDs are seen as 3490 devices).

Note: You do not need to define any specific esoterics (except in JES3) but you may want to define any number of esoterics to be referenced by SMC policies. For considerations for defining esoterics for ELS, refer to *Configuring and Managing SMC*.

Assigning MVS Device Numbers to RTDs

Use the HCD facility to assign MVS device numbers to these RTDs.

Tip: You must use the same unit addresses you determined for these transports for LIBGEN updates as described in "[Planning for Library-Attached Transports.](#)"

Setting the MVS Missing Interrupt Handler (MIH) Value

The VTSS's internal error recovery procedures requires the MVS missing-interrupt handler (MIH) value to be 20 minutes. You set this value by modifying the MIH parameter in SYS1.PARMLIB member IECIOSxx.

Note: When you are using VSM synchronous replication, set the MIH value to 45 minutes.

Configuring HSC

Before you configure VTCS, you must configure HSC if you are a new customer. If you are an existing customer, go to ["Reconfiguring a TapePlex."](#)

Your task list for configuring HSC is as follows:

- ["Creating the HSC CDS"](#)
- ["Defining Volumes to ELS"](#)
- ["Adding Definitions for ACF/VTAM Communications"](#)
- ["Defining the SYS1.PARMLIB Member SMFPRMxx"](#)
- ["Creating the HSC PARMLIB Member"](#)
- ["Defining Command Authority"](#)
- ["Updating HSM"](#)
- ["Creating and Cataloging the HSC Startup Procedure"](#)
- ["Performance Considerations"](#)

Note: For information on routing data sets to VSM, refer to *Configuring and Managing SMC*.

Creating the HSC CDS

For real tape configurations, creating the CDS consists of the following:

- ["Coding the LIBGEN Macros"](#)
- ["Running the SLICREAT Utility To Format the New CDS"](#)

For more information on the LIBGEN macros, see ["HSC LIBGEN Macros."](#)

For **only** Tapeless VSM configurations (no RTDs attached, can have VLEs attached), see ["Running the SLUADMIN CDSCREAT Utility to Create the CDS \(Tapeless VSM Only\)."](#)

Caution: It is critical that you control user access privileges to the CDS. See ["Limiting CDS Access Privileges"](#) for more information.

Coding the LIBGEN Macros

The LIBGEN macros specify the library configuration. You code the LIBGEN macros in a data set pointed to by the assembler JCL SYSIN DD statement. [Table 4-1](#) describes

the LIBGEN macros and shows their required order in the input data set. "[LIBGEN Macro Samples](#)" shows examples of LIBGEN input data sets. For more information on the LIBGEN utility, refer to the *ELS Command, Control Statement, and Utility Reference*. For more information on the LIBGEN macros, see "[HSC LIBGEN Macros](#)."

Caution: Oracle **strongly recommends** that you do not "pre-define" hardware through the LIBGEN process (for example, by defining transports that are not physically installed but will be in the future). Only use the LIBGEN process to define hardware that is actually installed. If you need to add hardware after performing an initial LIBGEN, use the techniques described in "[Reconfiguring a TapePlex](#)."

Table 4-1 LIBGEN Macros

Macro	Description
SLIRCVRY	Specifies HSC recovery values.
SLILIBRY	Specifies global settings. You can also define future ACSs with SLILIBRY.
SLIALIST	Specifies the list of ACSs.
SLIACS	To define the ACSs, create one SLIACS macro for each ACS listed in the SLIALIST macro.
SLISTATN	Specifies the list of the stations (LMU interfaces) that connect a host to an ACS. Specify one SLISTATN macro for each station entry listed in the SLIACS macro.
SLILSM	To define each LSM, create one SLILSM macro for each LSM listed in the SLIACS macro LSM, LSM2, LSM3, and LSM4 parameters.
SLIDLIST	To define the transport list for each LSM, create one SLIDLIST macro for each DRVELST parameter entry in the SLILSM macro. This SLIDLIST macro lists the SLIDRIVS macros for each host.
SLIDRIVS	Specifies the transport device addresses used by each host attached to an LSM.
SLIENDGN	Specifies the end of the LIBGEN macros.

To make this LIBGEN procedure real-world, create an input file step-by-step that ends up looking like the example file shown in "[One Host, Two ACSs, One 9310 LSM, One SL8500 Library Configuration](#)."

To create an example LIBGEN input file:

1. Create the SLIRCVRY macro:

```
LIBGEN01 SLIRCVRY TCHNIQE=ALL
```

In this macro, you specified the label for the generated CSECT, and specified all 3 copies of the CDS.

The TCHNIQE parameter on the SLIRCVRY LIBGEN macro determines how many CDS copies will be initialized by the SLICREAT program and whether or not journals will be initialized by SLICREAT. If journaling is specified by the TCHNIQE parameter, journals must be defined before HSC will successfully initialize.

2. Create the SLILIBRY macro:

```
SLILIBRY HOSTID=HSC1, X
SMF=255, X
```



```

COMPRFX=#,                X
ACSLIST=NAMEACS,         X
MAJNAME=STKALSQN,        X
SCRLABL=SL,              X
EJCTPAS=GOODAY,          X
CLNPRFX=CLN,             X

```

In this macro, you specified the following:

- Only host HSC1 accesses the configuration.
 - The SMF record type, here 255, which must match the type specified in ["Defining the SYS1.PARMLIB Member SMFPRMxx."](#)
 - # is the operator command prefix character.
 - NAMEACS is the assembler label of the SLIALIST macro.
 - STKALSQN is the ENQ/DEQ/QNAME used by host software for serialization.
 - Put Standard Labels on scratch volumes.
 - GOODAY is the eject password.
 - CLN is the cleaning cartridge prefix.
3. Create the SLIALIST macro:

```
NAMEACS  SLIALIST  ACS00,ACS01
```

In this macro, you specified the following:

- The assembler label is indeed NAMEACS, as specified in the SLILIBRY macro.
 - ACS00 and ACS01 are the two ACSs.
4. Create the SLIACS macro for ACS00:

```
ACS00    SLIACS  STATION=(ST000), X
LSM=(LSM0000)
```

In this macro, you specified the following:

- The assembler label is indeed ACS00, as specified in the SLIALIST macro.
 - The STATION parameter specifies ST000 as the assembler label of the SLISTATN macro.
 - The LSM parameter specifies LSM0000 as the assembler label of the (single) SLILSM macro for the 9310 LSM.
5. Create the succeeding macros referenced by the ACS00 SLIACS macro, starting with SLISTATN:

```
ST000    SLISTATN  ADDRESS=(0A0,0A1)
```

In this macro, you specified the following:

- The assembler label is indeed ST000, as specified in the SLIACS macro.
- The ADDRESS parameter specifies 0A0 and 0A1 as the device addresses of the two LMUs that control the 9310 LSM.

6. Create the SLILSM macro for the 9310 LSM:

```
LSM0000  SLILSM    DRIVE=(1,2,9,10),                X
          DRVELST=(P000001,P000002,P000009,P000010), X
          TYPE=9310,                                X
          DOOR=ECAP
```

In this macro, you specified the following:

- The assembler label is indeed LSM0000, as specified in the SLIACS macro.
- The DRVELST parameter specifies P0000001, P0000002, P0000009, and P000010 as the assembler labels of the SLIDLIST macros that define the transports in each transport panel.
- The LSM type is 9310.
- The access door has an Extended CAP.

7. Create the SLIDLIST and SLIDRVS macros for the 9310 LSM:

```
P000001  SLIDLIST  HOSTDRV=D0000010
*
D0000010 SLIDRIVS  ADDRESS=(410,411,412,413)
*
P000002  SLIDLIST  HOSTDRV=D0000020
*
D0000020 SLIDRIVS  ADDRESS=(414,415,416,417)
*
P000009  SLIDLIST  HOSTDRV=D0000090
*
D0000090 SLIDRIVS  ADDRESS=(510,511,512,513)
*
P000010  SLIDLIST  HOSTDRV=D000100
*
D000100  SLIDRIVS  ADDRESS=(, , )
```

In these macros, you specified the following:

- The assembler labels of the SLIDLIST macros that define the drive panels are indeed P0000001, P0000002, P0000009, and P0000010 as specified on the SLILSM macro.
- Each SLIDLIST HOSTDRV parameter specifies the assembler label of the corresponding SLIDRIVS macro.
- The assembler labels on the SLIDRVS macros match, and the ADDRESS parameters specify the unit addresses of each drive on each panel.

Note:

- These are the device addresses you determined in ["Planning for Library-Attached Transports."](#)
 - You did **not** specify drive addresses for panel P000010, so you can add drives there later if necessary.
-
-

Tip: This example shows a single host. For additional information on defining device addresses in a client/server environment, refer to *Configuring and Managing SMC*.

8. Create the SLIACS macro for ACS01:

```
ACS01    SLIACS FUTRLSM=(8500,16),           X
        LSM=(LSM0100,LSM0101,LSM0102,LSM0103)
```

Note: Before you configure HSC for the SL8500, verify that all components of the SL8500 are operational. HSC builds its library configuration from the information reported by the library. If SL8500 components are not operational, the library information may not be reported to HSC, and the HSC configuration of the SL8500 is incomplete.

To verify that all the components of the SL8500 are operational, use the SLConsole (either the panel on the SL8500 or a remote SLConsole). Select Tools, then select System Detail.

- All SL8500 components should be green. If the components are not green, and you have already configured the SL8500 to HSC, the missing components probably have not been configured.
 - If the drives are red, the drive configuration can be updated dynamically using the SET SLIDRIVS utility (refer to *ELS Command, Control Statement, and Utility Reference*) and the MODIFY CONFIG command. If only the drives are red, you can proceed with the HSC configuration.
 - The elevators (Elevator Folder) must be green. If the elevators are not green, do not configure the SL8500 to HSC. The elevators are the logical PTPs. Without PTPs, HSC will not know that the SL8500 rails are connected.
-

In this macro, you specified the following:

- The assembler label is indeed ACS01, as specified in the SLIALIST macro.
- The FUTRLSM parameter defines four SL8500s that will later be physically added to the configuration. HSC sees each rail in an SL8500 as an entire LSM. An SL8500 library contains four rails (LSMs), thus the lsmnum value must be in multiples of four, so in this example 16 LSMs equals 4 future SL8500s.
- The LSM parameter specifies the assembler labels of the 4 SLILSM macros for each SL8500 rail (starting with the first).

9. Create the SLILSM macro for the first SL8500 rail/LSM:

```
LSM0100  SLILSM PASTHRU=((0,M),(0,M),(0,M)),      X
        ADJACNT=(LSM0101,LSM0102,LSM0103),      X
        DRIVE=(1),                               X
        DRVELST=(P010001),                       X
        TYPE=8500,                                X
        DOOR=8500-1
```

In this macro, you specified the following:

- The assembler label is indeed LSM0100, as specified in the SLIACS macro.
- The PASTHRU parameter specifies:
 - 0, which is the value for SL8500 internal PTPs (elevators), repeated three times because the SL8500 has 4 LSMs and each LSM needs to pass thru to the other 3 LSMs.

- M designates this LSM as the master. In an SL8500, the lowest numbered LSM is always the master.
 - The ADJACNT parameter specifies the assembler labels of the other 3 rails (LSMs) in the SL8500.
 - The DRIVE parameter specifies only one transport panel, which is always the case for SL8500 rails/LSMs, and the DRVELST parameter specifies P010001 as the assembler label of the SLID LIST macro that defines the transports in the single drive panel.
 - The LSM type is 8500.
 - DOOR=SL8500-1 means that there is a single SL8500 CAP.
10. Create the SLIDLIST and SLIDRVS macros for the first SL8500 rail/LSM:

```
P010001 SLIDLIST HOSTDRV=(D0100010,D0100011)
*
D0100010 SLIDRVS ADDRESS=(,,,9400,,,,,9401,,,,,)
D0100011 SLIDRVS ADDRESS=(,,,9400,,,,,9401,,,,,)
```

In these macros, you specified the following:

- The assembler labels of the SLIDLIST macro that defines the single SL8500 drive panel is indeed P010001 as specified on the SLILSM macro.
- Each SLIDLIST HOSTDRV parameter specifies the assembler label of the corresponding SLIDRVS macro.
- The assembler labels on the SLIDRVS macros match, and the ADDRESS parameters specify the unit addresses of each drive on each panel.

Note:

- These are the device addresses you determined in "[Planning for Library-Attached Transports.](#)"
 - This example shows a drive panel for each of the two hosts. For more information on defining device addresses in a client/server environment, refer to *Configuring and Managing SMC*.
-
-

11. Repeat Step 8 through Step 10 for the other 3 SL8500 rails/LSMs.

The numbers will change but the basic coding does not except that these LSMs, in terms of PTP control, are Slave LSMs, which is always the case with SL8500s.

12. Create the SLIENDGN macro.

Running the SLICREAT Utility To Format the New CDS

After you code the LIBGEN macros as described in "[Coding the LIBGEN Macros](#)," assemble and link-edit the LIBGEN macros and run the SLICREAT utility to format the new CDS:

1. Assemble and link–edit the LIBGEN macros.

The following shows example LIBGEN assembler and linkage editor JCL. Update this example with your LIBGEN macros and run the job.

```
//*
```

```
//ASM      EXEC PGM=ASMA90
//SYSPRINT DD SYSOUT=*
//SYSTEM   DD SYSOUT=* (optional)
//SYSLIB   DD DSN=SYS1.MACLIB,DISP=SHR
//          DD DSN=SLS.SLSMAC,DISP=SHR
//SYSUT1   DD UNIT=SYSDA,SPACE=(CYL,(3,1))
//SYSLIN   DD DSN=&&OBJ,UNIT=SYSDA,
//          SPACE=(CYL,(1,1)),DISP=(,PASS)
//SYSIN    DD *
```

LIBGEN deck goes here

```
/*
//LKED     EXEC PGM=IEWL,PARM='LIST,XREF,RENT,REUS,REFR,RMODE=24',
//          COND=(0,NE)
//SYSPRINT DD SYSOUT=*
//SYSLMOD  DD DSN=SLS.your.hsc.linklib(lgename),DISP=SHR
//SYSUT1   DD UNIT=SYSDA,SPACE=(CYL,(3,1))
//SYSLIN   DD DSN=&&OBJ,DISP=(OLD,DELETE)
//*/
```

2. Run SLICREAT to format the New CDS.

The following shows example JCL for the SLICREAT program. This sample is also included in the HSC SAMPLIB as member JCLCRT.

```
//SLICREAT JOB (account),'programmer',CLASS=A
//CREATE EXEC PGM=SLICREAT, CDS CREATE MODULE
//          PARM='libgen-load-module-name',
//          REGION=6M
//*/
//STEPLIB DD DSN=your.hsc.linklib,DISP=SHR
//SYSPRINT DD SYSOUT=* MESSAGES
//*/
//*****
//*          LIBRARY PRIMARY CONTROL DATASET (CDS)
//*****
//SLSCNTL DD DSN=SLS.SLSCNTL, PRIMARY CDS
//          SPACE=(4096,s,,CONTIG,ROUND), REPLACE 's' WITH YOUR
//          DISP=(NEW,CATLG,DELETE), SPACE CALCULATIONS
//          UNIT=SYSDA
//*****
//*          LIBRARY SECONDARY CONTROL DATASET (CDS)
//*****
//SLSCNTL2 DD DSN=SLS.SLSCNTL2, SECONDARY CDS
//          SPACE=(4096,s,,CONTIG,ROUND), REPLACE 's' WITH YOUR
//          DISP=(NEW,CATLG,DELETE), SPACE CALCULATIONS
//          UNIT=SYSDA
//*****
//*          LIBRARY STANDBY CONTROL DATASET (CDS)
//*****
//SLSSTBY DD DSN=SLS.SLSSTBY, STANDBY CDS
//          SPACE=(4096,s,,CONTIG,ROUND), REPLACE 's' WITH YOUR
//          DISP=(NEW,CATLG,DELETE), SPACE CALCULATIONS
//          UNIT=SYSDA
//*****
```

In "CDS DASD Space Requirements," you determined the size of the CDS to support your ELS system.

Note:

- For an example of using SLICREAT to format journals, refer to the *ELS Legacy Interfaces Reference*.
 - Run SLICREAT with the highest level of the software that you are running on any LPAR.
-
-

Running the SLUADMIN CDSCREAT Utility to Create the CDS (Tapeless VSM Only)

For only new Tapeless VSM configurations (no RTDs attached, but you can have VLEs attached), run the SLUADMIN CDSCREAT utility to create the CDS as shown in the example below.

```
//CDSCREAT JOB (account), 'programmer', CLASS=A
//CREATE EXEC PGM=SLUADMIN,
//          REGION=6M
//*
//STEPLIB DD DSN=your.hsc.linklib, DISP=SHR
//SYSPRINT DD SYSOUT=*                MESSAGES
//*
//*****
//*          LIBRARY PRIMARY CONTROL DATASET (CDS)
//*****
//SLSCNTL DD DSN=SLS.SLSCNTL,          PRIMARY CDS
//          SPACE=(4096,s,,CONTIG,ROUND), REPLACE 's' WITH YOUR
//          DISP=(NEW,CATLG,DELETE),     SPACE CALCULATIONS
//          UNIT=SYSDA
//*****
//*          LIBRARY SECONDARY CONTROL DATASET (CDS)
//*****
//SLSCNTL2 DD DSN=SLS.SLSCNTL2,        SECONDARY CDS
//          SPACE=(4096,s,,CONTIG,ROUND), REPLACE 's' WITH YOUR
//          DISP=(NEW,CATLG,DELETE),     SPACE CALCULATIONS
//          UNIT=SYSDA
//*****
//*          LIBRARY STANDBY CONTROL DATASET (CDS)
//*****
//SLSSTBY DD DSN=SLS.SLSSTBY,          STANDBY CDS
//          SPACE=(4096,s,,CONTIG,ROUND), REPLACE 's' WITH YOUR
//          DISP=(NEW,CATLG,DELETE),     SPACE CALCULATIONS
//          UNIT=SYSDA
//*****
//SLSPRINT DD SYSOUT=*
//SLSIN DD *
CDSCREAT, +
HOSTID=(EC20,EC21), +
COMPRFX=4B, +
TCHNIQE=STANDBY
/*
```

Note: In "CDS DASD Space Requirements," you determined the size of the CDS to support your ELS system.

In this example:

- Hosts EC20 and EC21 are the SMF IDs of the hosts where the subsystem will execute.
- 4B is the hexadecimal representation of the command prefix character ('.')
- STANDBY indicates that three HSC CDS data sets are to be formatted and used by the system.

For details of all SLUADMIN CDSCREAT utility parameters, refer to the *ELS Commands, Control Statements, and Utility Reference*.

Note: In a tapeless configuration, you cannot issue commands or use parameters that reference library components.

Defining and Creating CDS Log Files

To define and create CDS log files:

1. Use the FMTLOG utility to pre-format the log files.

The size of the allocated log files depends on the number of transactions your system generates, as well as the frequency of performing the offload. StorageTek suggests you allocate and activate some test log files and determine the rate at which they fill before determining the size of the production log files.

Sample JCL:

```
//FMTLOG JOB (account),REGION=1024K
//S1 EXEC PGM=SLUADMIN,PARM=MIXED
//STEPLIB DD DSN=hlq.SEALINK,DISP=SHR
//SLSLOG1 DD DSN=hlq.CDSLOG1,DISP=(,CATLG,DELETE),
// UNIT=SYSDA,SPACE=(CYL,100)
//SLSLOG2 DD DSN=hlq.CDSLOG2,DISP=(,CATLG,DELETE),
// UNIT=SYSDA,SPACE=(CYL,100)
//SLSPRINT DD SYSOUT=*
//SLSIN DD *
FMTLOG
/*
```

Note: Data sets must be cataloged, so that they can be located by dynamic allocation. You can assign either one or two log files. If you assign two log files, the second is a copy of the first. There is no “automatic switching” between log files. If an error occurs on a file, it is disabled and logging continues on the other file.

2. Use the SET LOGFILE utility to activate the log files.

Sample JCL:

```
//SETLOG JOB (account),REGION=1024K
//S1 EXEC PGM=SLUADMIN,PARM=MIXED
//STEPLIB DD DSN=hlq.SEALINK,DISP=SHR
//SLSPRINT DD SYSOUT=*
//SLSIN DD *
SET LOGFILE(hlq.CDSLOG1,hlq.CDSLOG2)
/*
```

Running Multiple HSC Tapeplexes on One LPAR (MULT Mode)

ELS now supports the ability to run multiple instances of HSC on the same LPAR. The following restrictions apply when HSC is started in MULT mode:

- If any HSC on the LPAR specifies the MULT parameter, then all HSC subsystems on that LPAR must specify the MULT parameter.
- Two HSC subsystems on the same LPAR must represent separate tapeplexes, that is, each HSC must point to a different set of CDSs.
- ExPR started task is not supported on any LPAR where HSC is running in MULT mode.
- A new parameter, SSYS, is required on any SLUADMIN step running on an LPAR where HSC is in MULT mode.
- The SLIEXERS installation utility cannot be run on an LPAR where HSC is in MULT mode.
- Any application using the HSC Significant Event Notification (SEN) facility on an LPAR where HSC is running in MULT mode must be modified to pass a subsystem name parameter. Refer to the *ELS Programming Reference* for more detail on the SEN facility.

The following are recommendations for running multiple HSCs on one LPAR:

- Each HSC subsystem should specify a unique command prefix.
- Each HSC subsystem should specify MSGPRFX(YES) in its EXECPARM parameter. Using this parameter, along with a unique command prefix, identifies the HSC subsystem source for HSC console EXEC JCL parameters. Using this parameter, along with a unique command prefix, identifies the HSC subsystem source for HSC console messages.
- Each HSC subsystem should use a unique SMF record ID. This allows statistics to be collected for individual HSC subsystems.
- Each HSC subsystem should use a unique ENQUEUE name. This can be done by using the SLUADMIN SET MAJNAME utility. Specifying the same enqueue name for multiple CDSs can lead to performance issues.
- MULT mode and VTCS Coupling Facility:
 - If the VTCS Coupling Facility is used, StorageTek recommends that it should be used only on one HSC/VTCS system on an LPAR.
 - If the VTCS Coupling Facility is used on more than one HSC/VTCS subsystems within a single Sysplex, each system MUST specify a unique LOCKSTR value. Refer to the *ELS Command, Control Statement, and Utility Reference* for more information on the VTCS CONFIG utility.

Defining Volumes to ELS

With ELS 7.3 and above, you now use the HSC POOLPARM and VOLPARM statements to define all volumes and their pools, native Library Volumes, cleaning cartridges, MVCs, VTVs, and the HSC SET VOLPARM utility to load them. For more information, refer to the *ELS Command, Control Statement, and Utility Reference*.

Tip: Do you want to switch to POOLPARM/VOLPARM way but do not know how to do the conversion? The VOLPCONV utility helps to convert existing VOLDEF, SCRPDF, MVCDEF, CONFIG VTVVOL, and CONFIG MVCVOL statements to SET

VOLPARM statements. For more information, refer to the *ELS Command, Control Statement, and Utility Reference*.

In the following sections, you use the same basic procedure, with variations for the type of volume and its pool. But first, a discussion of the benefits of converting to POOLPARM/VOLPARM and some warnings about their use.

Why Should You Use POOLPARM/VOLPARM?

POOLPARM/VOLPARM provides the following advantages over the legacy method of defining volumes:

- Reduce system outages.
 - Using POOLPARM/VOLPARM allows you to dynamically delete VTV or MVC ranges from the CDS, without requiring a CDS MERGE process. Deleting VTV or MVC ranges dynamically requires CDS level G or above (CDSLEVEL=V62ABOVE) and also requires that none of the volumes in the deleted range be currently in use.
- Simplify maintenance of volume parameters.
 - All POOLPARM/VOLPARM parameters are maintained in the CDS.
 - Edits ensure that all volsers in your environment are defined only once and belong to only one pool.
 - Current contents of POOLPARM/VOLPARM is available by running the LIBGEN utility (refer to the *ELS Command, Control Statement, and Utility Reference* for details)
 - When the SET VOLPARM utility is run with APPLY=YES, all updates occur immediately on all hosts.
- Improved management of scratch subpool allocation.
 - POOLPARM/VOLPARM allows only non-subpool scratch volumes to be selected when no scratch subpool is specified. Non-subpool scratch volumes are those volumes defined before any pool definition.
 - Using POOLPARM/VOLPARM rejects mount requests for undefined subpool names from SMC.
 - POOLPARM/VOLPARM provides the ability to limit access to scratch subpools based on client host ID in a client/server environment.
- Improved reporting.
 - Your VOLRPT will automatically pick up subpool definitions, flag MVCs, and show volumes in DR test scratch subpools, as well as flagging undefined volumes.
 - HSC D VOLUME DETAIL shows pool residency.
- Pre-validation.
 - You can use the SLUADMIN utility to pre-validate your POOLPARM/VOLPARM statements before activating them.
 - You can use the SMC POLICY VALIDATE in the SMC subsystem or SMCUSIM utility to ensure that all SMC POLICY statements specify valid subpools.
- POOLPARM/VOLPARM supports volume segregation in the Concurrent Disaster Recovery Test utility.

- Using POOLPARAM/VOLPARAM automatically segregates DR test scratch subpools and MVC pools from production. When POOLPARAM/VOLPARAM is in effect, DR test systems can use only volumes defined in DRTEST pools.
- In a DR test environment, you can scratch volumes defined as part of DRTEST scratch subpools.
- Using POOLPARAM/VOLPARAM supports volume segregation in a Cross TapePlex Replication environment.
 - Defining POOLPARAM with TYPE=EXTERNAL and OWNRPLEX name ensures that the receiving TapePlex cannot use a VTV as scratch and can receive a VTV only from the specified OWNRPLEX
 - Changing a POOLPARAM from TYPE=EXTERNAL to TYPE=SCRATCH allows the receiving TapePlex to scratch VTVs received from another TapePlex in case of an actual disaster.
- A utility is provided to assist in converting existing volume definitions to POOLPARAM (for more information, refer to the *ELS Command, Control Statement, and Utility Reference*).
 - Input includes existing VOLDEF, SCRPFDEF, MVCPOOL, and VTCS CONFIG definitions
 - Note that the conversion utility cannot convert "wildcard" characters in VOLATTR statements. You must manually convert wildcard specifications to volume ranges.

Some Things to Consider with POOLPARAM/VOLPARAM

A few things to remember when using POOLPARAM/VOLPARAM:

- VTCS CONFIG MVCVOL and VTVVOL definitions are no longer honored.
- HSC control commands SCRPFDEF, VOLDEF and MVCDEF are no longer honored.
- Your pool definitions are the same on all hosts. You can, however, use the HOSTID parameter to limit access to scratch subpools.
- Mounts for invalid subpools will be rejected.
- Mounts with no specified subpool can be satisfied only by non-subpool volumes.
- You can disable POOLPARAM/VOLPARAM by running the SET VOLPARAM utility with an empty SLSPARM input file. Once POOLPARAM/VOLPARAM is disabled, the previous version of POOLPARAM/VOLPARAM remains in effect for MVC and VTV definitions until you run VTCS CONFIG, and the previous POOLPARAM/VOLPARAM POOL definitions remain in effect until you run the VOLDEF, SCRPFDEF and MVCDEF commands on the HSC host(s).
- If you run CDRT with POOLPARAM/VOLPARAM in effect you must have DRTEST pools defined or scratch mounts will not be satisfied, and no MVCs can be used.
- SET VOLPARAM execution time increases exponentially as more volume ranges are added. To facilitate efficient processing of VOLPARAM statements, combine volume ranges.

For example, the following ranges can be combined into one VOLPARAM statement using the SET VOLPARAM JOIN utility.

```
VOLPARAM VOLSER(A00000-A00999) MED(VIRTUAL) RECT(VIRTUAL)
VOLPARAM VOLSER(A01000-A01999) MED(VIRTUAL) RECT(VIRTUAL)
VOLPARAM VOLSER(A02000-A02999) MED(VIRTUAL) RECT(VIRTUAL)
```

For more information on the SET VOLPARM JOIN utility, refer to the *ELS Command, Control Statement, and Utility Reference*.

Defining Library Volumes

To define Library Volumes:

1. Create a POOLPARM statement to define volume scratch pools.

For example, to define a scratch pool for use by hosts MVS1 and MVS2:

```
POOLPARM NAME (SYS19840P1) TYPE (SCRATCH) HOSTID (MVS1 , MVS2) LABEL (SL)
```

2. Create a VOLPARM statement to define the volumes.

For example, to define a range of 9840C volumes:

```
VOLPARM VOLSER (C1000-C1999) MEDIA (STK1R) RECTECH (STK1RC)
```

3. Ensure that you define sufficient cleaning cartridges.

Go to ["Defining Cleaning Cartridges."](#)

Note: Once you install higher density (or encryption-capable) drives for existing media, note that if you eject and re-enter an existing cartridge, HSC sets the density to "unknown" whenever the ACS contains drives other than the lowest density for the media, and sets the encryption status to "unknown" whenever the ACS contains encrypting drives for the media. This means that in the allocation process, SMC attempts to allocate the specific volume to the highest density drives and to encrypting drives, whenever the volume is flagged as "unknown."

To manage this situation, do one of the following:

- You can use VOLPARMs for the volume to indicate its actual density (encryption status), and then, once it is scratched, reset the VOLPARM to the desired recording technique.
 - You can use the vaulted volume feature to manage the ejected cartridge as a non-library volume, using LCM. When you use this feature, the HSC CDS records the recording technique at the time of the eject and restores it at the time of the enter.
-
-

Defining Cleaning Cartridges

To define cleaning cartridges:

1. Create a POOLPARM statement to define the cleaning cartridge pool.

For example, to define the 9840 cleaning cartridge Pool with cleaning parameters:

```
POOLPARM TYPE (CLEAN) MAXCLEAN (25)
```

2. Create a VOLPARM statement to define the cleaning cartridges.

For example, to define a range of 9840 cleaning cartridges:

```
VOLPARM VOLSER (CLN300-CLN310) MEDIA (STK1U)
```

3. Do you need to define MVCs?

If so, go to ["Defining MVCs."](#) Otherwise, go to ["Validating and Applying the Volume Definitions."](#)

Defining MVCs

Note: POOLPARAM/VOLPARAM is the preferred method for defining MVCs for ELS 7.3 but you can still use CONFIG to define MVCs. For more information, refer to the *ELS Legacy Interfaces Reference*. Also note that if you have existing volsers defined using CONFIG, you must run DECOMP against your CONFIG deck, remove the CONFIG definitions, and rerun CONFIG before you create POOLPARAM/VOLPARAM definitions.

To define MVCs:

1. Create a POOLPARAM statement to define the MVC pool.

For example, to define the T10000 MVC Pool with reclaim parameters:

```
POOLPARAM NAME(SYS1MVCT1) TYPE(MVC) MVCFREE(40) MAXMVC(4) THRESH(60) START(70)
```

2. Create a VOLPARAM statement to define the MVCs.

For example, to define a range of T10000 Full volumes to be encrypted:

```
VOLPARAM VOLSER(T2000-T2999) MEDIA(T10000T1) RECTECH(T1AE)
```

3. Define VTVs.

See ["Defining VTVs."](#)

Defining VTVs

Note: POOLPARAM/VOLPARAM is the preferred method for defining VTVs for ELS 7.3 but you can still use CONFIG to define MVCs. For more information, refer to the *ELS Legacy Interfaces Reference*. Also note that if you have existing volsers defined using CONFIG, you must run DECOMP against your CONFIG deck, remove the CONFIG definitions, and rerun CONFIG before you create POOLPARAM/VOLPARAM definitions.

To define VTVs:

1. Create a POOLPARAM statement to define the volume scratch pool.

For example, to define a VTV scratch pool for use by hosts MVS1 and MVS2:

```
POOLPARAM NAME(SYS1VTV1) TYPE(SCRATCH) HOSTID(MVS1, MVS2) LABEL(SL)
```

2. Create a VOLPARAM statement to define the VTVs.

For example, to define a range of VTVs initialized as scratch:

```
VOLPARM VOLSER(V5000-V5999)MEDIA(VIRTUAL) INITSCR
```

Validating and Applying the Volume Definitions

1. Run SET VOLPARM to validate the POOLPARM/VOLPARM statements.

```
SET VOLPARM APPLY(NO)
```

APPLY(NO) validates the statements without loading them. If you like the results, go to Step 2. Otherwise, rework your volume definitions, then go to Step 2.

2. Run SET VOLPARM to load the POOLPARM/VOLPARM statements.

```
SET VOLPARM APPLY(YES)
```

3. Physically enter the cartridges into the ACS.

For more information, refer to *Managing HSC and VTCS*.

Adding Definitions for ACF/VTAM Communications

To optimize performance, Oracle recommends that you set the HSC COMMPATH METHOD parameter to TCP, which does not place a performance burden on HSC, as shown in the example in "[Creating the HSC PARMLIB Member](#)."

If you set the HSC COMMPATH METHOD to TCP, then specify the SYSTCPD DD statement in your HSC START procedure to define TCP/IP options for the HSC job, as shown in the example in "[Sample HSC Startup Procedure](#)" on page 4-23. The SYSTCPD DD statement identifies the data set used to obtain parameters defined by the IBM TCPIP.DATA configuration data set. Refer to the *IBM TCP/IP Customization and Administration Guide* for more information.

An alternative to setting the HSC COMMPATH METHOD to TCP is to set the HSC COMMPATH METHOD to VTAM. VTAM is another HSC host-to-host communications method that does not place a performance burden on HSC.

If you set COMMPATH METHOD to VTAM, you must define the following to VTAM:

- APPL
- CDRSC
- CDRM (if not using existing CDRMs)
- LOGMODE table entry for SNASVCMG (contained in the IBM-supplied logon mode table).
- LOGMODE table entry for SLSSVCMG. This can be defined the same as SNASVCMG. If not defined, the default logmode entry is used.

A sample APPL statement is contained in the HSC SAMPLIB member HSCAPPL, shown below, with the HSC application program minor node (APPLID) as APHSC1:

```
HSCAPPL  VBUILD TYPE=APPL
APHSC1  APPL  APPC=YES,
          AUTOSES=1,          +
          DMINWNL=1,         +
          DMINWNR=1,         +
          DESESLIM=2,        +
          EAS=10
```

where:

APPC=YES

must be coded because HSC uses VTAM LU 6.2 services.

AUTOSES=1

defines the number of contention-winner sessions VTAM is to establish automatically with the first CNOS request.

DMINWNL=1

defines the minimum number of parallel sessions with the local LU as the contention-winner. HSC only requires and uses one local contention-winner session.

DMINWNR=1

defines the minimum number of parallel sessions with the remote LU as the contention-winner. HSC only requires and uses one remote contention-winner session.

DSESLIM=2

defines the maximum number of sessions allowed between the local LU and a remote LU. This should be 2 because the HSC only establishes two sessions between each HSC: one local contention-winner session and one remote contention-winner session.

EAS=10

sets an estimated number of concurrent sessions this APPLID will have with other LUs. Refer to the IBM ACF/VTAM manuals for definitions and explanations.

Defining the SYS1.PARMLIB Member SMFPRMxx

HSC can produce SMF record subtypes for HSC and VTCS events. To produce these record subtypes, you must add two statements to your SMF parameters in SYS1.PARMLIB member SMFPRMxx to specify the following:

- HSC subsystem for which records are produced
- Recording interval in seconds
- SMF record subtypes. The record subtypes must be specified as a list (*subtype1, subtype2,...subtypen*), as a range (*subtype1-subtypen*), or as a combination (*subtype1, subtype2-subtypen*). A range must be specified using a dash; a colon is invalid for a range.

Note: The TYPE parameter in the SMFPRMxx member must match the SLILIBRY SMF= parameter specified in "[Coding the LIBGEN Macros](#)."

Tip: Oracle recommends that you specify that your system produces the HSC SMF record subtypes 1 through 8 and 10, 11, 13 through 21, and 25 through 33 as shown in the following example. For more information on SMF record layouts, refer to the *ELS Programming Reference*.

The following shows sample statements that produce record subtypes 1 through 8 and 10, 11, 13 through 21, and 25 through 33 at 15 minute intervals for HSC subsystem SLS0.

```
SUBSYS(SLS0, INTERVAL(001500), TYPE(255))
SUBPARM(SLS0(SUBTYPE, (1-8,10,11,13-21,25-33)))
```

Creating the HSC PARMLIB Member

The HSC PARMLIB member is where you specify any commands that you want executed at HSC startup as shown in the example that follows.

```
CDSDEF DSN1 (HSC1.PRIM) , DSN2 (HSC1.BKUP) , DSN3 (HSC1.STDBY)
MGMTDEF DSN (HSC.PARMS)
COMMPATH HOSTID (HST1) TCPPATH (HST1NAME METHOD (TCP)
COMMPATH HOSTID (HST2) TCPPATH (HST2NAME METHOD (TCP)
CAPP 9,01:00:00,AUTO
MNTD AUTOCLN (ON) EJCTAUTO (ON)
```

In this example:

CDSDEF

The CDSDEF control statement in the HSC PARMLIB determines how many CDS copies are allocated by the HSC. These definitions are invoked at HSC initialization and remain set until HSC termination. The definitions cannot be altered without HSC shutdown and restart.

Caution: The number of CDS copies implied by the TCHNIQE parameter must be the same as the number of CDS copies implied by the CDSDEF control statement otherwise the HSC will not initialize.

- If the number of CDS copies defined on the CDSDEF control statement is more than are specified by the TCHNIQE parameter, do the following:

Stop the HSC on all hosts.

Run a SLUADMIN SET TCHNIQE to increase the number of copies to match the number implied by the CDSDEF control statement.

Restart HSC.

- If the number of CDS copies defined on the CDSDEF control statement is less than are specified by the TCHNIQE parameter, do the following:

Stop the HSC on all hosts.

Run a SLUADMIN SET TCHNIQE to decrease the number of copies to match the number implied by the CDSDEF control statement.

Restart HSC.

The control data set names are recorded in the Database Heartbeat (DHB) record. During HSC initialization, the data set names in the DHB are compared with the data set names specified on the CDSDEF control statement in the HSC PARMLIB.

If a data set name specified on the CDSDEF statement does not match any of the control data set names recorded in the DHB, that control data set is disabled. If all the specified control data sets are disabled, the HSC will not initialize. The specific assignment of enabled control data sets to the primary, secondary, and standby control data set(s) is made based on the control data set assignments recorded in the Database Heartbeat record.

HSC.PARMS

is the data set that contains your system’s MGMTclas and STORclas statements.

**COMMPATH HOSTID(HST1) TCPPATH(HST1NAME) METHOD(TCP)
COMMPATH HOSTID(HST2) TCPPATH(HST2NAME) METHOD(TCP)**

specifies that TCP is the host-to-host communications method between HSC hosts HST1 and HST2. Oracle recommends that you specify TCP to allow even sharing of resources in a multi-host environment

CAPP 9,01:00:00,AUTO

specifies CAP preferences.

MNTD AUTOCLN(ON) EJCTAUTO(ON)


specifies automatic drive cleaning and automatic ejection of spent cleaning cartridges.

EXECPARM Statement

The EXECPARM statement provides an alternative method of specifying GTF Eid and Fid parameters.

Syntax

Figure 4–1 EXECPARM Statement



Parameters

MSGPRFX

optionally, specifies control of whether the command prefix is to precede WTO or WTOR messages to the operator.

Yes

indicates that the command prefix is to display preceding WTO or WTOR messages to the operator.

No

indicates that the command prefix is not to display preceding WTO or WTOR messages to the operator.

Eid

optionally, *gtfeid* specifies a GTF event ID.

PARM='Eid(*user-specified-event-id*)' parameter is valid for use in the HSC initialization procedure as an alternative method of specifying the GTF event ID.

Fid

optionally, *gtffid* specifies a GTF format ID.

PARM='Fid(*user-specified-format-id*)' parameter is valid for use in the HSC initialization procedure as an alternative method of specifying the GTF format ID.

HOSTID

optionally, *host-id* specifies the system ID associated with the request to process the EXECParM control statement.

Note: If the *hostid* specified does not match the host executing the command, the command is ignored and no message is issued.

Defining Command Authority

If HSC user exit SLSUX15 sets a return code of UX15CHKA, the exit issues a command authorization request to the system security product. The following shows an example of RACF profile and permissions commands to give user SAM15 update access to all HSC/VTCS commands and user JOE13 query access to all HSC/VTCS commands. Note that you can no longer specify the .VT prefix for VTCS commands, which is no longer passed to SLSUX15. Refer to the *ELS Programming Reference* for a sample of how to define command security based on the originating host ID of the command.

```
*****
* Define a profile in the OPERCMDS class for all HSC and VTCS commands *
*****
RDEFINE OPERCMDS subsystem.* UACC (NONE)
*****
***** Allow user JOE13 access to all query HSC/VTCS commands
*****
*****PERMIT subsystem.command CLASS(OPERCMDS) ID(JOE13) ACCESS (QUERY)
*****
***** Allow user SAM15 update access to all HSC/VTCS commands
*****
*****PERMIT subsystem.command CLASS(OPERCMDS) ID(SAM15) ACCESS (UPDATE)
*****
```

Updating HSM

HSM users that have mixed devices that were "logically" defined as the same type of device, such as 3490E but are "physically" different, such as T9940, virtual (VTD), or 9490 must set the following parameter in HSM:

```
SETSYS RECYCLEINPUTDEALLOCFREQUENCY(MIGRATION(1))
```

By setting this parameter, when HSM is "recycling," it will deallocate the input drive after it processes each input tape. This is required where the tapes being recycled are "physically" mixed as described above.

If you do not set this parameter, it is possible that you could allocate a 9490 transport for the first tape, then if the second tape was virtual (VTV) or STK2P, the job would fail due to media incompatibility. That is, you could not physically mount the second tape (virtual or STK2P media) on the 9490 drive that had been allocated for the first tape.

Creating and Cataloging the HSC Startup Procedure

You start HSC by entering the MVS START command that invokes the HSC startup procedure, which resides in the cataloged procedure library of the host system. For more information, see:

- ["Running Multiple HSC Tapeplexes on One LPAR \(MULT Mode\)"](#)
- ["EXEC Statement"](#)
- ["Sample HSC Startup Procedure"](#)
- ["EXECPARM Statement"](#) for an alternative method of specifying GTF Eid and Fid parameters.

Note: If you previously controlled which hosts started VTCS by referencing or not referencing the VTCS software in your JCL concatenation, this capability is provided in ELS by using the NOVTCS parameter on the HSC startup PARM. For more information, see ["EXEC Statement."](#)

Running Multiple HSC Tapeplexes on One LPAR (MULT Mode)

ELS supports the ability to run multiple instances of HSC on the same LPAR. The following restrictions apply when HSC is started in MULT mode:

- If any HSC on the LPAR specifies the MULT parameter, then **all** HSC subsystems on that LPAR must specify the MULT parameter.
- Two HSC subsystems on the same LPAR must represent separate tapeplexes, that is, each HSC must point to a different set of CDSs.
- The ExPR started task is not supported on any LPAR where HSC is running in MULT mode.
- A new parameter, SSYS, is required on any SLUADMIN step running on an LPAR where HSC is in MULT mode.
- The SLIEXERS installation utility cannot be run on an LPAR where HSC is in MULT mode.
- Any application using the HSC Significant Event Notification (SEN) facility on an LPAR where HSC is running in MULT mode must be modified to pass a subsystem name parameter. Refer to the *ELS Programming Reference* for more detail on the SEN facility.

The following are recommendations for running multiple HSCs on one LPAR:

- Each HSC subsystem should specify a unique command prefix.
- Each HSC subsystem should specify MSGPRFX(YES) in its EXECPARM or EXEC JCL parameters. Using this parameter, along with a unique command prefix, identifies the HSC subsystem source for HSC console EXEC JCL parameters. Using this parameter, along with a unique command prefix, identifies the HSC subsystem source for HSC console messages.
- Each HSC subsystem should use a unique SMF record ID. This allows statistics to be collected for individual HSC subsystems.
- Each HSC subsystem should use a unique ENQUEUE name. This can be done by using the SLUADMIN SET MAJNAME utility. Specifying the same enqueue name for multiple CDSs can lead to performance issues.
- MULT mode and VTCS Coupling Facility.
 - If the VTCS Coupling Facility is used, Sun StorageTek recommends that it should be used only on one HSC/VTCS system on an LPAR.


- If the VTCS Coupling Facility is used on more than one HSC/VTCS subsystems within a single Sysplex, each system **must** specify a unique LOCKSTR value. Refer to the *ELS Command, Control Statement, and Utility Reference* for more information on the VTCS CONFIG utility.

EXEC Statement

The EXEC statement specifies HSC startup options.

Syntax

Figure 4-2 EXEC Statement



Parameters

PARM=

defines the list of parameters passed to the HSC initialization routine.

Note: If you enter more than one of the following parameters, you must separate them with a blank space (e.g., BASE SSYS(subsystem) RESET).

BASE

specifies that the HSC initialize and execute at the base service level.

SSYS

specifies that HSC initialization search for the subsystem name specified. If the name is not found or is not a valid name, the subsystem terminates. subsystem must be a 1- to 4-character name or problems can occur when initializing the HSC.

This parameter permits you to symbolically specify the subsystem and retain the same startup procedure whether starting the HSC before or after JES.

INIT

specifies that only preinitialization of the HSC occurs.

Note: If PARM=INIT is specified, the HSC subsystem is only initialized. It is still necessary to issue a HSC Start command to start the HSC. If any other parameters are specified with INIT (except for SSYS), they are ignored.

COLD

specifies that any permanent in-memory data structures previously allocated by HSC are reallocated and reinitialized.

On the first startup of the HSC after an IPL, this option is meaningless. If the HSC has been brought up previously for this IPL, use of this option results in the loss of a system linkage index for Program Call (PC) instructions. There are a limited number of system linkage indexes. Once exhausted, they can only be restored by IPLing. If COLD is not specified, the linkage index used previously by the HSC is reused.

This parameter should be used only when absolutely necessary. (The installation instructions for some HSC maintenance may direct you to perform a COLD start.)

Note: You do not need to include the COLD parameter when you are initializing a HSC that is at a different release level than the HSC that was previously running on a host. When an initializing HSC detects a release level difference, it performs an automatic internal cold start.

RESET

specifies that all subsystem status flags in the MVS Subsystem Communications Vector Table (SSCVT) for the HSC are unconditionally reset. Use of this option may correct a situation in which the HSC was terminated abnormally without resetting the status flags; for example, if the HSC was terminated with the MVS FORCE command.

These messages are possible symptoms of this situation at HSC startup, when a display of active jobs indicates that the subsystem is not active:

- ... ACS subsystem CCCC is ACTIVE
- or
- ... ACS subsystem CCCC is TERMINATING
- or
- ... ACS subsystem CCCC is INITIALIZING

This parameter should only be used in extreme situations and may not correct all error conditions. Contact Oracle Software Support before using this parameter.

Eid

xxxx is 1 to 4 hex characters specifying the GTF event ID used for the duration of this subsystem. "E" is the abbreviation for this parameter. The default Eid value is **E086**.

Fid

xx is 1 to 2 hex characters specifying the GTF format ID used for the duration of this subsystem. "F" is the abbreviation for this parameter. The default Fid value is **17**.

Note: See ["EXECPARM Statement"](#) for an alternate method of specifying GTF, EID, and FID parameters.

Member

For MVS, *xx* is the suffix of the SLSSYS*xx* member in SYS1.PARMLIB, or an SLSSYS*xx* DD statement in the startup procedure used as the automatic commands (PARMLIB control statements) data set. "M" is the abbreviation for this parameter.

Dialog

specifies that messages can be displayed on the operator console and/or written to the system log. This option can be used to further restrict where messages are displayed based on the ROUTCDE. These messages indicate that the HSC is waiting for an active task to complete before the HSC terminates.

For more information on Dialog, see "OPTION Command and Control Statement" in the *ELS Command, Control Statement, and Utility Reference*.

If Dialog is specified, one of the options **must** be selected. There is no default. The options for Dialog include:

Off

specifies that you do not want active task termination messages displayed on the operator console or written to the system log.

Both

specifies that messages are displayed on the operator console and written to the system log.

Console

specifies that messages are displayed on the operator console only.

Log

specifies that messages are written to the system log only.

NOVTCS

specifies that VTCS should not initialize on a host even though VTCS data is present in the CDS.

MULT

Optionally, specifies that multiple HSCs can be run on the same LPAR. Note that if one HSC on an LPAR specifies the MULT parameter, **all** HSCs running on that LPAR must specify the MULT parameter. See ["Running Multiple HSC Tapeplexes on One LPAR \(MULT Mode\)"](#) for a list of restrictions when using this parameter.

Sample HSC Startup Procedure

The following shows HSC SAMPLIB member JCLPROC, which is a sample HSC startup procedure.

```
//SLS0      PROC  PROG=SLSBINIT,PRM=' `
//*
//IEFPROC   EXEC  PGM=&PROG,TIME=1440,
//  PARM='&PRM E(E086) F(23) M(00) SSYS(SLS0)',REGION=6M
//*
//STEPLIB  DD   DSN=SEA700.SEALINK,DISP=SHR
//*
//SLSSYS00 DD   DSN=SLS.PARMS,DISP=SHR
```

```
//SYSTCPD DD DSN=ddd.eee.fff(anyname) /* Optional TCPIP parameters) */
```

In this example:

- In the example PROC shown above, SLS0 is used as both the name of the startup procedure (line 1) and as the subsystem name defined in the SSYS parameter "SSYS(SLS0)" in line 4. This is just an example, you can code any values that work for your site but note the following:
 - HSC subsystem names must be 4 characters in length and must be defined in your IEFSSNxx member.
 - If the startup procedure name and the subsystem name are identical, the SSYS parameter is not required in the startup procedure. If the startup procedure name is **not** identical to the subsystem name, then the SSYS parameter must point the startup procedure to the required HSC subsystem.
 - TIME=1440 or TIME=NOLIMIT must be specified to ensure that HSC does not time out and terminate.
 - Oracle recommends that you run HSC/VTCS with a region size of at least 6 MB except if you are running utilities or commands that manipulate manifest files, in which case you need the maximum region size your system will allow.

DD Statements for the HSC Startup Procedure

SLSSYSxx

statement that defines the data set containing the HSC PARMLIB. In the example procedure above, *xx* is replaced by the suffix "00". The SLSSYS00 DD statement matches the M(00) declaration and points to the PARMLIB member "00" that contains your start-up parameters.

SLSUEXIT

statement that defines the data set containing HSC user exits.

Specifying TCPIP Options The SYSTCPD DD statement can be used to specify TCP/IP options for the HSC job.

This DD statement identifies the data set used to obtain parameters defined by the IBM TCPIP.DATA configuration data set. Refer to the IBM *TCP/IP Customization and Administration Guide* for more information.

Performance Considerations

This section describes performance considerations.

Define Dispatching Priority for ELS Software

Depending on your requirements, set these software components so the amount of resources available to them is as follows:

- HSC/VTCS — greater than batch and any started task or application requiring quick access to tape but less than an online system
- SMC/HTTP server— greater than the HSC, or if the HSC is inactive, greater than batch and any started task or application requiring quick access to tape but less than an online system

Your requirements may dictate that these products be given access to greater or lesser resources than Oracle recommends here, but you must determine this for your environment. The information above is meant only to serve as a guideline.

During initialization, the HSC uses SYSEVENT TRANSWAP to make its address space non-swappable. The HSC address space cannot be swapped out once this has been accomplished.

Reduce Pass-Thrus

The number of pass-thrus required to mount, dismount, and replace cartridges in LSMs can impact library performance. In a large or busy ACS, this impact may be significant especially during periods of heavy mount activity. There are three types of pass-thrus:

- unavoidable
- unnecessary
- scheduled.

Unavoidable Pass-Thrus

The HSC attempts to minimize the number of pass-thrus required; however, depending upon available tape transports and locations of cartridges, pass-thrus often cannot be avoided. By running Activities Reports on a regular basis and examining the results, you can see that mounts for different LSMs take longer than mounts for the same LSM.

Unnecessary Pass-Thrus

Scratch mounts, dismounts, enters, and ejects that require pass-thrus are unnecessary and should be avoided. These types of activities divert the robotics from productive work especially during periods of peak activity.

If scratch subpools are properly defined and managed, scratch cartridges normally are not involved in pass-thrus. Scratch mounts are the same as specific mounts as far as pass-thrus that cannot be avoided.

The effect of unnecessary pass-thrus is not obvious in the Activities Report. You must compare their number with average mount times to see the effect on performance.

Scheduled Pass-Thrus

Using the Scratch Redistribution utility to balance scratch cartridges involves moving cartridges to various LSMs to achieve scratch balancing. This type of activity involves heavy pass-thru usage. If you must balance scratch volumes across your library, schedule such activity during off peak hours. This approach will make sure that the pass-thru activity involved does not directly interfere with mounts and dismounts for high-priority production.

You can use LCM to schedule pass-thrus. Refer to the *LCM User's Guide* for more information.

Methods to Reduce Pass-Thru Activity

There are various methods to reduce pass-thru activity:

Set MNTD Float to ON

The Float option of the MNTD command specifies whether the HSC is to select a new home cell location when it dismounts a volume that requires a pass-thru when it is dismounted.

When MNTD Float is set to ON, cartridges are not passed back to their original LSM. The cartridges are assigned new cells in the LSM where they are dismounted. This action eliminates most unnecessary pass-thrus.

Provide adequate free cells

Setting MNTD Float to ON can be overridden if there are no free cells in the dismounting LSM. Dismounted cartridges are passed to other LSMs to find a new home cell.

Use the Display Lsm command to determine the number of free cells in each LSM. Use MOVE or EJECT to create free cells if they are needed.

Eject through the CAP closest to the cartridge

If you eject a cartridge through the CAP of the LSM where it resides, no pass-thrus are required.

If you eject a cartridge without specifying a CAPid, the cartridge is ejected through the highest priority CAP that is not busy. This type of activity may cause one or more unnecessary, nonproductive pass-thrus.

The recommended way to accomplish ejects without affecting performance with pass-thrus, is to use multiple CAP option on the EJECT utility. By specifying multiple CAPs (i.e., CAP(00:00:00,00:00:01,00:00:02)), the desired effect (i.e., no pass-thrus) is achieved.

Redistribute cartridges during off-peak times

You can use the MOVE command and utility to move cartridges within an LSM or between LSMs. The Scratch Redistribution utility can be used to move scratch cartridges between LSMs until an equilibrium is reached. Each inter-LSM movement of cartridges causes pass-thrus which delay robot movement in mounting a cartridge.

Depending on the number of cartridges to be redistributed, you may prefer to schedule moves and scratch redistribution during periods of low data center activity. The redistribution runs faster and performance is not affected during off peak times.

Reduce Operator Intervention

Although the ACS runs mostly unattended, situations occur where operator intervention is required. Excessive and unnecessary operator intervention impacts library performance. There are specific ways in which operator intervention can be reduced. These include:

Set SMC ALLOCDef DEFer(ON)

If a keep is issued for a cartridge while it is being mounted, the HSC issues a message indicating that the robot cannot dismount the volume. An operator must unload the tape transport and then reply "R" to the message.

This situation is caused by programs allocating cartridges through JCL, not opening the data set on the cartridge, and terminating before the mount completes. It causes unnecessary mounts, extends dismount time, and delays the availability of cartridges and transports.

If a message requiring operator intervention occurs frequently, you may want to set ALLOCDef DEFer(ON) so ACS mounts are deferred until the data set is opened.

Alternatively, you can code User Exit 09 for JES2 or User Exit 11 for JES3 with SETUP processing to selectively defer a subset of ACS mounts.

Set CAP Preference

In a multiple-LSM ACS, the CAPPref command establishes an ordered list of CAPs to use should the operator or HSC start an activity that requires a CAP without specifying a CAPid.

To enter or eject cartridges faster, set CAPPref such that the CAP closest to the cartridge racks is preferred. This minimizes operator travel distance.

In large ACS configurations, of five or more LSMs, consider setting CAPPref such that a CAP in the middle of the ACS has the highest priority. This may make the operators walk further but it reduces the number of pass-thrus should the default (highest priority) CAPid be used for ejecting cartridges.

Prefetch Enters

Nonlibrary cartridges that are mounted on library transports are delayed while the operator fetches and enters them. This is a common occurrence for HSC sites with all transports attached to the library.

If you or your scheduling system can predict which nonlibrary cartridges will be mounted in the library before a mount message appears, your operator can improve performance by entering those cartridges in advance.

Avoid Crashing Test Systems

Library attached hosts own library resources, including CAPs, tape drives, and cartridges. If a host fails, another host must clean up the resources held by the failing host. This delays mounts and dismounts on the recovering host while recovery takes place.

You should attempt to shut down the HSC properly before IPLing a system. This is especially important for test systems that may be restarted several times a day.

Use Unit Affinity Separation

Unit affinity separation can improve library performance by eliminating the need for operators to enter or eject tape cartridges when volumes are requested with unit affinity. Without unit affinity separation, a tape transport is selected based on the location of the first volume.

Refer to the *SMC Configuration and Administration Guide* for information about the ALLOCDef command SEPLvl parameter, which specifies the exclusion level at which affinity chains are separated.

Reduce Tape Transport Contention

Balanced use of library tape transports results in better robotic and system performance. In a multiple LSM library, you want the workload to be spread evenly among the robots rather than overloading one robot while the others are idle. Within each LSM, you want mounts to be evenly distributed among transports rather than having the robot wait for a cartridge to be rewound so it can mount the next cartridge on the same tape transport.

Tape transport contention can be reduced by:

- ensuring scratch cartridge balance
- managing multi-host tape transports
- avoiding dedicating tape transports
- using two tape transports for multi-volume files.

Reduce Scheduling Contention

Effective scheduling can increase library performance. Controlling the following scheduling related areas can help significantly in further increasing library performance:

- strive for a balanced workload
- schedule nonproductive library activities during low-demand times.

Balancing the Workload

If you experience higher than expected average mount times but at the same time experience an improvement in production throughput, it could be an indication that your system is periodically flooded with work rather than having a balanced workload.

For example, all of your production jobs may be getting submitted at the beginning of a shift so that the library robots are overworked for the first few hours of the shift. Then, the robots may remain idle for the remainder of the shift. If you are using a scheduling software package and it releases jobs every hour on the hour, there may be tremendous tape transport contention for the first few minutes of each hour while the transports are unused for the remainder of each hour.

These situations tend to elevate average mount times; however, as long as the work is performed on time, there is probably no need to change conditions. However, if the work is not getting performed on schedule, you can improve performance by balancing your production workload.

Scheduling Nonproductive Activities During Low-Demand Times

There are several library utilities, which, though very essential, severely impact the library's ability to mount and dismount cartridges. The following utilities should not be run when high-priority production jobs are pending:

- AUDIt
- EJECT
- Initialize Cartridge
- MOVE
- Scratch Redistribution
- Scratch Update.

These utilities should be scheduled during quiet periods so they do not contend with mounts and dismounts. Running these activities during quiet periods also makes sure that the tasks complete faster. In the case of mass enters or ejects, operator's time can also be optimized.

Reconfiguring a TapePlex

"[Configuring HSC](#)" explained how to configure HSC, which includes performing a LIBGEN, which is how you initially define hardware to HSC. This chapter is for existing customers who want to add hardware to their systems either on release boundaries or between release boundaries.

The idea is that adding hardware to an existing system is usually easier, because you do not have to redefine each and every piece of hardware, and also because adding hardware is an easier path.

For example, to add transports, you do not have to stop HSC, recreate your entire LIBGEN, and so forth. You can just use SET SLIDRIVS and dynamic reconfiguration (using MODIFY CONFIG) to make it happen.

When Do You Need to Reconfigure Your TapePlex?

Whenever the hardware components that make up your tapeplex change, you need to perform a reconfiguration task to synchronize the HSC CDS with the physical tape hardware. Some reasons for reconfiguring a tapeplex include:

- Consolidating two data centers
- Splitting a single data center into two tapeplexes
- Adding, changing or removing physical tape hardware, including drives, LSMs, and ACSs.
- Adding, changing or removing VTSSs.
- Adding, changing or removing VAULT information.

Depending on the existing configuration and the type of change being done, you will need to use one of the following tools/procedures:

- HSC MERGEcds utility
- HSC dynamic reconfiguration process
- VTCS CONFIG utility
- HSC SET VOLPARM utility
- HSC SET VAULT and/or SET VAULTVOL utility
- HSC AUDIT utility
- HSC SET CDKLOCK utility

The MERGEcds Utility

The MERGEcds utility is used to migrate information from an existing CDS into a new, formatted CDS. The MERGEcds utility is normally used for the following types of changes:

- Consolidating or splitting data centers
- Adding LSMs or ACSs to a configuration when FUTRLSM or FUTRACS was not specified or was insufficient for the new configuration.
- Deleting LSMs or ACSs from a configuration
- Changing panel types in a non-Streamline library
- Deleting or renaming VTSSs
- Changing the number of vaulted volumes in a CDS where the reserved number of vaulted volumes is too small (or, when the number of vaulted volumes is larger than needed and you want to decrease the size of the CDS)
- Changing the size of the CDKLOCK subfile. The MERGEcds utility does not copy the content of the old CDKLOCK subfile to the new CDS.

Basic MERGEcds Procedure

The following sections describe how to use the MERGEcds utility to add an ACS to an existing HSC. This procedure is the basis for using MERGEcds for the examples in the other section, such as merging data centers.

Create a New, Empty CDS

Note: This procedure can be run at any time before the actual MERGEcds operation.

To create a new, empty CDS:

1. Run the HSC LIBGEN utility to create LIBGEN macros that reflect your current real hardware configuration (and, optionally, your VOLPARM and vaulted volume data). For example:

```
//SLSLIBGN JOB (ACCT), 'LIBGEN', NOTIFY=&SYSUID
//SLSLIBGN EXEC PGM=SLUADMIN, PARM='MIXED'
//STEPLIB DD DSN=hlq.SEALINK, DISP=SHR
/* NOTE: CDS DD STATEMENTS NOT NEEDED IF HSC ACTIVE
//SLSPRINT DD SYSOUT=*
/* NOTE: SET VOLPARM, SET VAULTVOL, SET VAULT ARE
/* OUTPUT IN THE DATASET BELOW
//SLSSET DD DSN=hlq.SLSSET, DISP=(NEW, CATLG, DELETE),
// SPACE=(CYL, 1)
/* NOTE: SET VOLPARM POOLPARM/VOLPARM CONTROL STATEMENTS
/* ARE IN THE DATA SET BELOW.
//SLSPARM DD DSN=hlq.SLSPARM, DISP=(NEW, CATLG, DELETE),
// SPACE=(CYL, 1)
/* NOTE: LIBGEN MACROS ARE OUTPUT IN THE DATA SET
/* BELOW
//SLSLIBGN DD DSN=hlq.SLSLIBGN, DISP=(NEW, CATLG, DELETE),
// SPACE=(CYL, 1)
//SLSIN DD *
LIBGEN
```

- //
2. Update the SLSLIBGN output file to add the new ACS to your configuration.
See ["Configuring HSC"](#) and samples in ["LIBGEN Macro Samples"](#) for more information on updating LIBGEN macros.
 3. If necessary, update your VOLPARM and POOLPARM statements to reflect new volume ranges.
 4. If necessary, update the SET VAULTVOL, SET VAULT, and SET CDKLOCK statements to reflect your desired new configuration.
 5. Assemble the updated LIBGEN module. For example:

```
//SLSASSML JOB (ACCT), 'ASSM LIBGEN', NOTIFY=&SYSUID
//ASM      EXEC PGM=ASMA90,
//          PARM='XREF (SHORT), TERM, TEST, RENT'
//SYSLIB   DD DSN=hlq.SEAMAC, DISP=SHR
//          DD DSN=SYS1.MACLIB, DISP=SHR
//SYSPRINT DD SYSOUT=*
//SYSTEM  DD SYSOUT=*
//*      NOTE: UPDATED LIBGEN MACROS ARE IN THE DATA SET
//*      BELOW
//SYSIN    DD DSN=hlq.LIBGEN, DISP=SHR
//SYSLIN   DD DSN=&&TMPOBJ, DISP=(, PASS, DELETE),
//          UNIT=SYSDA, SPACE=(CYL, (1, 1)),
//          DCB=(LRECL=80, BLKSIZE=3040, RECFM=FB)
//SYSUT1   DD DSN=&SYSUT1, UNIT=VIO, SPACE=(CYL, (1, 1))
//*
//LKED     EXEC PGM=IEWL, COND=(0, LT, ASM), REGION=6M,
//          PARM='LIST, LET, XREF, TEST, RENT'
//SYSLIN   DD DSN=&&TMPOBJ, DISP=(OLD, DELETE)
//SYSPRINT DD SYSOUT=*
//SYSLMOD  DD DSN=hlq.LIBGNLIB(LGYMMDD), DISP=SHR
//SYSUT1   DD UNIT=VIO, SPACE=(CYL, (1, 1))
//
```

6. If your configuration includes VSM data, run the VTCS DECOMP utility to create the current source for your VSM configuration. For example:

```
//SLSDECOM JOB (ACCT), 'VTCS DECOMP', NOTIFY=&SYSUID
//SLSDECOM EXEC PGM=SLUADMIN, PARM='MIXED'
//STEPLIB DD DSN=hlq.SEALINK, DISP=SHR
//*      NOTE: CDS DD STATEMENTS NOT NEEDED IF HSC ACTIVE
//SLSPRINT DD SYSOUT=*
//*      NOTE: CONFIG CONTROL CARDS ARE OUTPUT IN THE DATA SET
//*      BELOW
//CONFIG   DD DSN=hlq.CONFIG, DISP=(NEW, CATLG, DELETE),
//          SPACE=(CYL, 1)
//SLSIN    DD *
//          DECOMP FLATDD(CONFIG)
//
```

7. Run the SLICREAT utility to determine the number of blocks needed for real tape processing, using the output of the LIBGEN assembly. For example:

```
//SLICREA1 JOB (ACCT), 'SLICREAT SIZE', NOTIFY=&SYSUID
//SLICREA1 EXEC PGM=SLICREAT, PARM=LGYMMDD, REGION=6M
//STEPLIB DD DSN=hlq.SEALINK, DISP=SHR
//          DD DSN=hlq.LIBGNLIB, DISP=SHR
//SYSPRINT DD SYSOUT=*
```

```
//
```

8. If your configuration includes VSM data, vaulted volume data, and/or VOLPARM control statements, see ["CDS DASD Space Requirements"](#) to calculate the additional space needed.
9. Run the SLICREAT utility to create the new CDS, using the your block calculation for the new CDS size. For example:

```
//SLICREA2 JOB (ACCT), 'SLICREAT CDS', NOTIFY=&SYSUID
//SLICREA1 EXEC PGM=SLICREAT, PARM=LGYMMDD, REGION=6M
//STEPLIB DD DSN=hlq.SEALINK, DISP=SHR
// DD DSN=hlq.LIBGNLIB, DISP=SHR
//SLSCNTL DD DSN=hlq.DBASEPRM.NEW, DISP=(NEW, CATLG, DELETE) ,
// UNIT=SYSDA, SPACE=(4096, xxxx, , CONTIG)
//SLSCNTL2 DD DSN=hlq.DBASESEC.NEW, DISP=(NEW, CATLG, DELETE) ,
// UNIT=SYSDA, SPACE=(4096, xxxx, , CONTIG)
//SLSSTBY DD DSN=hlq.DBASESBY.NEW, DISP=(NEW, CATLG, DELETE) ,
// UNIT=SYSDA, SPACE=(4096, xxxx, , CONTIG)
//SYSPRINT DD SYSOUT=*
//
```

10. Run the SLUADMIN SET utility to add the VOLPARM definitions, vaulted volume count, and vault information and and CDKLOCK to the new CDS. For example:

```
//SLSSETCD JOB (ACCT), 'SET CDS', NOTIFY=&SYSUID
//SLSLIBGN EXEC PGM=SLUADMIN, PARM='MIXED'
//STEPLIB DD DSN=hlq.SEALINK, DISP=SHR
//SLSCNTL DD DSN=hlq.DBASEPRM.NEW, DISP=SHR
//SLSCNTL2 DD DSN=hlq.DBASESEC.NEW, DISP=SHR
//SLSSTBY DD DSN=hlq.DBASESBY.NEW, DISP=SHR
//SLSPRINT DD SYSOUT=*
/* NOTE: ORIGINAL OR MODIFIED POOLPARM/VOLPARM CONTROL
/* STATEMENTS ARE IN THE DATA SET BELOW
//SLSPARM DD DSN=hlq.SLSPARM, DISP=SHR
/* NOTE: ORIGINAL OR MODIFIED SET COMMANDS FOR
/* VOLPARM, VAULTVOL, AND VAULT ARE IN THE DATA SET
/* BELOW
//SLSIN DD DSN=hlq.SLSSET, DISP=SHR
//
```

11. Run the VTCS CONFIG utility to create the CDS VSM data areas. For example:

```
//SLSCONFG JOB (ACCT), 'VTCS CONFIG', NOTIFY=&SYSUID
//SLSCONFG EXEC PGM=SLUADMIN, PARM='MIXED'
//STEPLIB DD DSN=hlq.SEALINK, DISP=SHR
//SLSCNTL DD DSN=hlq.DBASEPRM.NEW, DISP=SHR
//SLSCNTL2 DD DSN=hlq.DBASESEC.NEW, DISP=SHR
//SLSSTBY DD DSN=hlq.DBASESBY.NEW, DISP=SHR
//SLSPRINT DD SYSOUT=*
/* NOTE: ORIGINAL OR MODIFIED VTCS CONFIG CONTROL
/* STATEMENTS ARE IN THE DATA SET BELOW
//SLSIN DD DSN=hlq.CONFIG, DISP=SHR
//
```

Prepare to Run MERGEcds

To prepare to run MERGEcds:

1. Issue the Display CDS command to determine the primary CDS.
2. Quiesce tape activity on all HSC and client hosts.
3. Stop HSC on all active hosts.
4. Run the SLUADMIN BACKUP utility to back up your current production CDS.
For example:

```
//SLSBKUP JOB (ACCT), 'CDS BKUP', NOTIFY=&SYSUID
//SLSBKUP EXEC PGM=SLUADMIN, PARM= 'MIXED'
//STEPLIB DD DSN=hlq.SEALINK, DISP=SHR
//SLSCNTL DD DSN=hlq.DBASEPRM, DISP=SHR
//SLSCNTL2 DD DSN=hlq.DBASESEC, DISP=SHR
//SLSSTBY DD DSN=hlq.DBASESBY, DISP=SHR
//SLSBKUP DD DSN=hlq.CDSBKUP,
//          DISP= (NEW, CATLG, DELETE) ,
//          UNIT=SYSDA, SPACE= (CYL, xxx) ,
//          DCB= (LRECL=4096, BLKSIZE=4096, RECFM=F)
//SLSPRINT DD SYSOUT=*
//SLSIN DD *
BACKUP
//
```

5. Run the SLUADMIN RESTORE utility to restore your production CDS to a new name. For example:

```
//SLSRESTR JOB (ACCT), 'CDS RESTORE', NOTIFY=&SYSUID
//SLSRESTR EXEC PGM=SLUADMIN, PARM= 'MIXED'
//STEPLIB DD DSN=hlq.SEALINK, DISP=SHR
//SLSCNTL DD DSN=hlq.DBASEPRM.OLD, DISP=SHR
//SLSCNTL2 DD DSN=hlq.DBASESEC.OLD, DISP=SHR
//SLSSTBY DD DSN=hlq.DBASESBY.OLD, DISP=SHR
//SLSBKUP DD DSN=hlq.CDSBKUP, DISP=SHR
//SLSPRINT DD SYSOUT=*
//SLSIN DD *
RESTORE
//
```

Note: Instead of backing up and restoring the production CDS, you can rename the primary CDS (from your Display CDS command) to a new name, and use that CDS as input to the MERGEcds.

6. Rename your new CDS(s) to the production names.

Using the names in the above JCL samples, rename hlq.DBASEPRM.NEW to hlq.DBASEPRM, rename hlq.DBASESEC.NEW to hlq.DBASESEC, and rename hlq.DBASESBY.NEW to hlq.DBASESBY.

7. Modify your HSC PARMLIB CAPPREF statement to reflect CAPs in the new ACS.

For example, to set priority to 9 (highest), mode to automatic to CAP 00, LSM 00 in new ACS 01:

```
CAPP 9,01:00:00,AUTO
```

Run the MERGEcds Utility

To run the MERGEcds utility:

1. Start HSC at BASE level, specifying the new CDS(s) in the CDSDEF statement.

2. Run the SLUADMIN MERGECDS utility.

For example, MERGEcDS VALIDATE:

```
//SLSMERGV JOB (ACCT), 'CDS MERGE VALIDATE', NOTIFY=&SYSUID
//SLSMERGV EXEC PGM=SLUADMIN, PARM='MIXED'
//STEPLIB DD DSN=hlq.SEALINK, DISP=SHR
//SLSFCNTL DD DSN=hlq.DBASEPRM.OLD, DISP=SHR
//SLSPRINT DD SYSOUT=*
//SLSIN DD *
MERGECDS ALL VALIDATE
//
```

For example, actual merge:

```
//SLSMERGE JOB (ACCT), 'CDS MERGE', NOTIFY=&SYSUID
//SLSMERGE EXEC PGM=SLUADMIN, PARM='MIXED'
//STEPLIB DD DSN=hlq.SEALINK, DISP=SHR
//SLSFCNTL DD DSN=hlq.DBASEPRM.OLD, DISP=SHR
//SLSPRINT DD SYSOUT=*
//SLSIN DD *
MERGECDS ALL
//
```

Ensure the Merge was Successful

To ensure that the merge was successful:

1. Review the merge output to ensure that all the CDS pieces were successfully merged.
2. Run the SLUADMIN BACKUP utility to back up the new CDS.
3. Bring HSC to FULL service level, issuing the command:

```
SRVLEV FULL
```

4. Connect the new ACS to the HSC, issuing the command:

```
Vary ACS acs-id Online
```

5. Modify the new LSMs online, issuing the command:

```
F LSM acs-id:lsm-id Online
```

6. Resume normal tape activity.
7. If the merge output showed duplicate volumes or errant or in-transit volumes, use the SLUADMIN AUDIT utility to audit the home cells of these cartridges.

In the following examples, the basic steps for executing the MERGEcDS procedure are identical to the process described above. These examples illustrate additional steps and MERGE control cards that are needed for other scenarios.

Consolidating Two Data Centers

In consolidating two data centers, you assume that each data center has its own CDS.

To merge the two CDSs:

1. Create a new CDS that reflects the library, virtual, and vaulted volume sections of both input CDSs as described in ["Create a New, Empty CDS."](#)

2. Next, put data from the first CDS into the new CDS as described in ["Prepare to Run MERGEcds"](#) and ["Run the MERGEcds Utility."](#)
3. To merge data from the second CDS, you need to run an additional MERGEcds step. For example, actual merge:

```
//SLSMERG2 JOB (ACCT), 'CDS MERGE 2', NOTIFY=&SYSUID
//SLSMERG2 EXEC PGM=SLUADMIN, PARM='MIXED'
//STEPLIB DD DSN=hlq.SEALINK, DISP=SHR
//SLSFCNTL DD DSN=hlq.DBASEPRM.CDS2, DISP=SHR
//SLSFCTL2 DD DSN=hlq.DBASESEC.CDS2, DISP=SHR
//SLSFSBY DD DSN=hlq.DBASESBY.CDS2, DISP=SHR
//SLSPRINT DD SYSOUT=*
//SLSIN DD *
MERGECDs
/* NOTE: In THE FOLLOWING EXAMPLE, DATA CENTER 1
/* HAD 4 ACSS AND DATA CENTER 2 HAD 2 ACSS
/* DATA CENTER 1 HAD A VAULT NAMED VAULT1, AND
/* DATA CENTER 2 ALSO HAD A VAULT NAMED VAULT1.
//SLSMERGE DD *
MERGE FACS(00) TACS(04)
MERGE FACS(01) TACS(05)
ALLVIRT /* NO VTSS NAME CONFLICT */
MERGE FVALT(VAULT1) TVALT(VAULT5)
//
```

4. Ensure that the merge was successful as described in ["Ensure the Merge was Successful."](#)

Combining Two Separate ACSs into One ACS

The process of combining two ACSs into a single ACS requires a MERGEcds to change the definition of LSMs in the ACS being merged.

To combine two separate ACSs into one:

1. Create a new CDS that reflects the existing library hardware plus the new combined ACS, and the virtual and vaulted volume data as described in ["Create a New, Empty CDS."](#)
2. Next, put data into the new CDS as described in ["Prepare to Run MERGEcds"](#) and ["Prepare to Run MERGEcds,"](#) using the following sample JCL.

Merging two ACSs:

```
//SLSMERG3 JOB (ACCT), 'CDS MERGE 3', NOTIFY=&SYSUID
//SLSMERG3 EXEC PGM=SLUADMIN, PARM='MIXED'
//STEPLIB DD DSN=hlq.SEALINK, DISP=SHR
//SLSFCNTL DD DSN=hlq.DBASEPRM.OLD, DISP=SHR
//SLSPRINT DD SYSOUT=*
//SLSIN DD *
MERGECDs
/* NOTE: In THE FOLLOWING EXAMPLE, ACSS 02 (4 LSMS)
/* AND 03 (4 LSMS) ARE COMBINED INTO A SINGLE ACS.
//SLSMERGE DD *
MERGE FLSM(03:00) TFSM(02:04)
MERGE FLSM(03:01) TFSM(02:05)
MERGE FLSM(03:02) TFSM(02:06)
MERGE FLSM(03:03) TFSM(02:07)
ALLVIRT
ALLVALT
```

//

3. Check for warning messages if panel types in either ACS have changed.
If any panel types are changed, run the SLUADMIN AUDIT utility to audit the contents of the changed panels.
4. Ensure that the merge was successful as described in "[Ensure the Merge was Successful.](#)"

Combining Two Separate Streamline Libraries

SL8500 PTP support allows multiple, separate SL8500s to be combined into a single ACS by running the MERGEcds utility. To minimize losing information about cartridges and to reduce operational downtime, follow the procedures described in this section.

Two configurations are shown below. In the first configuration (see [Figure 5-1](#) and [Figure 5-2](#)), the ACSs to be merged are numbered from right to left (as you face them from the CAP side). In the second, they are numbered from left to right.

It is assumed that the resulting ACS uses the lower or lowest ACSid, which is ACS0. The placement and numbering of the pre-merge ACSs determines how volume data is mapped from the original ACSs' LSMs into the merged ACSs LSMs.

In an ACS with multiple SL8500s, the LSMs are always numbered from right to left, so volumes in the left most SL8500 must be in the four highest numbered LSMs. Conversely volumes in the right most LSM must be in the four lowest numbered LSMs.

Once the individual ACSs are physically merged into a single ACS, the MERGEcds utility is used to map each LSM of the pre-merge ACSs into an LSM in the merged, single ACS. The physical location of the volumes within the merged ACS is identical to what it was prior to the merge but their location designator, their LSMid, has probably changed.

Configurations for ACSs Numbered Right to Left

[Figure 5-1](#) and [Figure 5-2](#) illustrate merging three ACSs numbered right to left, as you face them from the CAP side.

Figure 5-1 Existing Right to Left Configuration - Three Separate ACSs

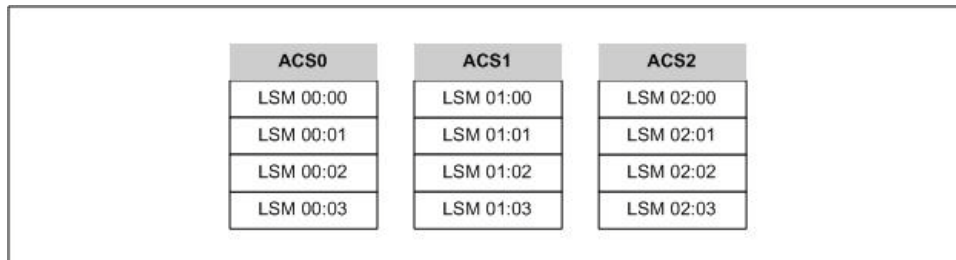


Figure 5–2 Desired Configuration - Single ACS

ACS0					
LSM 00:08	P T P	LSM 00:04	P T P	LSM 00:00	
LSM 00:09		LSM 00:05		LSM 00:01	
LSM 00:0A		LSM 00:06		LSM 00:02	
LSM 00:0B		LSM 00:07		LSM 00:03	

Configurations for ACSs Numbered Left to Right

Figure 5–3 and Figure 5–4 illustrate merging three ACSs numbered left to right, as you face them from the CAP side.

Figure 5–3 Existing Left to Right Configuration - Three Separate ACSs

ACS0	ACS1	ACS2
LSM 00:00	LSM 01:00	LSM 02:00
LSM 00:01	LSM 01:01	LSM 02:01
LSM 00:02	LSM 01:02	LSM 02:02
LSM 00:03	LSM 01:03	LSM 02:03

Figure 5–4 Desired Configuration - Single ACS

ACS0					
LSM 00:08	P T P	LSM 00:04	P T P	LSM 00:00	
LSM 00:09		LSM 00:05		LSM 00:01	
LSM 00:0A		LSM 00:06		LSM 00:02	
LSM 00:0B		LSM 00:07		LSM 00:03	

To combine two Separate Streamline libraries:

1. Create a new CDS for the two libraries as described in ["Create a New, Empty CDS."](#)
2. Enter data into the new CDS as described in ["Prepare to Run MERGEcdfs"](#) and ["Run the MERGEcdfs Utility."](#) The following JCL examples show control cards for this example, first using right to left numbering of ACSs and next using left to right numbering of ACSs.

SLSMERGE Statements - ACSs Numbered Right to Left When you merge LSMs numbered right to left, the LSM numbers of the "from" LSM and the "to" LSM in the MERGE statements both increase regularly as you move from right to left merging individual LSMs. Figure 5–5 shows SLSMERGE statements with ACSs numbered right to left.

Figure 5-5 SLSMERGE Statements Right to Left

```

MERGE FLSM(01:00) TLSM(00:04)
MERGE FLSM(01:01) TLSM(00:05)
MERGE FLSM(01:02) TLSM(00:06)
MERGE FLSM(01:03) TLSM(00:07)
MERGE FLSM(02:00) TLSM(00:08)
MERGE FLSM(02:01) TLSM(00:09)
MERGE FLSM(02:02) TLSM(00:0A)
MERGE FLSM(02:03) TLSM(00:0B)
MERGE FLSM(00:00) TLSM(00:00)
MERGE FLSM(00:01) TLSM(00:01)
MERGE FLSM(00:02) TLSM(00:02)
MERGE FLSM(00:03) TLSM(00:03)
    
```

SLSMERGE Statements - ACSs Numbered Left to Right When you merge LSMs numbered left to right, the LSM numbers of the original ACSs are increasing in the opposite direction from those in the merged ACS, where LSMs are number right to left. Therefore, as you merge LSMs moving right to left, the LSM numbers of the "to" LSMs increase regularly but the LSM numbers of the "from" LSMs decrease (in blocks of four increasing numbers). [Figure 5-6](#) shows SLSMERGE statements with ACSs numbered left to right.

Figure 5-6 SLSMERGE Statements Left to Right

```

MERGE FLSM(00:00) TLSM(00:08)
MERGE FLSM(00:01) TLSM(00:09)
MERGE FLSM(00:02) TLSM(00:0A)
MERGE FLSM(00:03) TLSM(00:0B)
MERGE FLSM(01:00) TLSM(00:04)
MERGE FLSM(01:01) TLSM(00:05)
MERGE FLSM(01:02) TLSM(00:06)
MERGE FLSM(01:03) TLSM(00:07)
MERGE FLSM(02:00) TLSM(00:00)
MERGE FLSM(02:01) TLSM(00:01)
MERGE FLSM(02:02) TLSM(00:02)
MERGE FLSM(02:03) TLSM(00:03)
    
```

These updates must occur first

These updates must occur last

Note: When you merge LSMs numbered from left to right, the order the SLSMERGE statements are processed is important. In the example below, the "to" LSMs for the pre-existing ACS 2 (02:00-02:03) become the "from" LSMs 00:00-00:03. You must first update ACS 0, LSMs 00:00-00:03 to be "from" LSMs 00:08-00:0B. Otherwise, the volumes from both "from" LSMs 00:00-00:03 and 02:00-02:03 will be assigned to 00:08-00:0B.

Ensure that the merge was successful as described in ["Ensure the Merge was Successful."](#)

Removing PTPs and Splitting an ACS Using MERGEcdfs

Oracle recommends you locate the new ACS to the left side of the split, which ensures that HSC assigns higher-numbered LSMs to the new ACS and avoids renumbering the LSMs that remain in the existing LSM.

If you create a new ACS to the right side of the split, all volume addresses must be remapped using MERGEcdfs. The procedures below do not discuss creating the new ACS to the right of the split.

To remove PTPs and split an ACS using MERGEcdfs:

1. Create a new CDS as described in ["Create a New, Empty CDS."](#)
2. Enter data into the new CDS as described in ["Prepare to Run MERGEcdfs"](#) and ["Run the MERGEcdfs Utility"](#) using the following sample JCL.

```
//SLSMERG7 JOB (ACCT), 'CDS MERGE 7', NOTIFY=&SYSUID
//SLSMERG7 EXEC PGM=SLUADMIN, PARM='MIXED'
//STEPLIB DD DSN=hlq.SEALINK, DISP=SHR
//SLSFCNTL DD DSN=hlq.DBASEPRM.OLD, DISP=SHR
//SLSPRINT DD SYSOUT=*
//SLSIN DD *
MERGECDS
//* NOTE: SLSMERGE STATEMENTS, ACS NUMBERED RIGHT
//* TO LEFT
//SLSMERGE DD *
MERGE FLSM(00:08) TLSM(01:00)
MERGE FLSM(00:09) TLSM(01:01)
MERGE FLSM(00:0A) TLSM(01:02)
MERGE FLSM(00:0B) TLSM(01:03)
MERGE FLSM(00:0C) TLSM(01:04)
MERGE FLSM(00:0D) TLSM(01:05)
MERGE FLSM(00:0E) TLSM(01:06)
MERGE FLSM(00:0F) TLSM(01:07)
//
```

3. Ensure that the merge was successful as described in ["Ensure the Merge was Successful."](#)

Renaming a VTSS

A VTSS rename can be done only by using the MERGEcdfs utility, which allows VTV and MVC data in the CDS to reflect the new VTSS name.

To rename a VTSS using MERGEcdfs:

1. Create a new CDS as described in ["Create a New, Empty CDS."](#)
2. Next, put data into the new CDS as described in ["Prepare to Run MERGEcdfs"](#) and ["Run the MERGEcdfs Utility,"](#) using the following sample JCL:

```
//SLSMERG4 JOB (ACCT), 'CDS MERGE 4', NOTIFY=&SYSUID
//SLSMERG4 EXEC PGM=SLUADMIN, PARM='MIXED'
//STEPLIB DD DSN=hlq.SEALINK, DISP=SHR
//SLSFCNTL DD DSN=hlq.DBASEPRM.OLD, DISP=SHR
//SLSPRINT DD SYSOUT=*
//SLSIN DD *
MERGECDS
//SLSMERGE DD *
ALLREAL
MERGE FVTSS(VTSS1) TVTSS(VTSSA)
```

```
    ALLVALT
//
```

3. Ensure that the merge was successful as described in ["Ensure the Merge was Successful."](#)

Removing Vaulted Volume Information

To remove vaulted volume information from the CDS, you must run the Merge process:

1. Create a new CDS, removing the SET VAULTVOL and SET VAULT control cards from the LIBGEN SLSSET output file as described in ["Create a New, Empty CDS."](#)
2. Enter data into the new CDS as described in ["Prepare to Run MERGEcds"](#) and ["Run the MERGEcds Utility,"](#) using the following sample JCL:

```
//SLSMERG5 JOB (ACCT), 'CDS MERGE 5', NOTIFY=&SYSUID
//SLSMERG5 EXEC PGM=SLUADMIN, PARM='MIXED'
//STEPLIB DD DSN=hlq.SEALINK, DISP=SHR
//SLSFCNTL DD DSN=hlq.DBASEPRM.OLD, DISP=SHR
//SLSPRINT DD SYSOUT=*
//SLSIN DD *
MERGECDs
//SLSMERGE DD *
    ALLREAL
    ALLVIRT
    NOVALT
//
```

3. Ensure that the merge was successful as described in ["Ensure the Merge was Successful."](#)

Reconfiguring a Tapeplex Without MERGEcds

The following topics address tapeplex reconfiguration without using MERGEcds.

- ["Using Dynamic Reconfiguration"](#)
- ["Using Audit Instead of MERGEcds for Streamline Libraries"](#)

Using Dynamic Reconfiguration

Dynamic reconfiguration, which you do using MODIFY CONFIG, is a key tool for NCO, the idea being that MODIFY CONFIG lets you change your hardware configuration without having to recycle HSC. Specifically, MODIFY CONFIG does the following:

- Modifies a CAP or an LSM online or offline to all hosts, independent of the issuing host.
- Starts dynamic hardware reconfiguration for a SL3000 or a SL8500 library.
- Adds, deletes, or updates an SL3000 or SL8500 ACS definition in the CDS that was preallocated using the SLLIBRY FUTRACS parameter.

There are some guidelines for dynamic reconfiguration, however, see below.

Caution:

- **Do not** run any CDS-related commands or utilities while making NCO changes.
 - Back up the CDS before and after you perform a hardware reconfiguration.
 - For any active HSC that accesses the CDS being modified by dynamic reconfiguration, **all ACSs and LSMs must be varied online** to maintain configuration integrity across hosts. An inactive HSC can be started after dynamic reconfiguration has completed on another host but once an HSC becomes active, it must be at the full service level, and all ACSs and LSMs must be varied online.
 - In a multi-host environment, Oracle recommends that **all active hosts run HSC 6.2-level or above software**. Any hosts not at 6.2 level or above **must** be brought down before initiating dynamic hardware reconfiguration and brought back up after the hardware reconfiguration is complete. The non-6.2 hosts recognize the hardware changes at this time.
-

Adding LSMs

Adding LSMs without FUTRLSM defined requires running the MERGEcds utility as described in "[Basic MERGEcds Procedure](#)." If you have defined FUTRLSM, you can use Dynamic Reconfiguration to add a new LSM as follows:

1. Use dynamic reconfiguration to bring the new LSM online:

```
MODIFY CONFIG
```

2. If necessary, add transports and volumes. For more information, see:
 - "[Defining Volumes to ELS](#)"
 - "[Adding Transports](#)"
3. If necessary, populate LSM by any of the following:
 - Entering volumes through CAPs.
 - Opening the access door, manually populating storage cells, then doing an audit of all cells.
 - Let PTPs move volumes from attached LSMs.

Dynamically Reconfiguring Streamline Library Transports

This section tells how to do the following:

- "[Adding Transports](#)"
- "[Removing Transports](#)"
- "[Replacing Transports](#)"

Adding Transports

Note:

- If you enter new cartridges into the library, update your SET VOLPARM definitions to include the new volumes. Then review SMC POLICY statements to determine if you want to add or modify policies if you are adding a new type of transport to your environment.
 - If your new transports will be RTDs, then update your VTCS CONFIG.
 - If you are adding a new RTD device type, update the STORCLAS definitions for the storage classes that will be directed to the new device type.
-
-

This section provides a procedure to add drives to an Streamline Library and HSC configuration. In this case, the Streamline Library panel already contains operational transports, and the HSC remains active during this time, so NCO is used to add the transports.

To dynamically add transports to a Streamline Library:

1. Install the new transports and power them on.
2. Use the SET SLIDRIVS utility to add the MVS drive address and drive location to the CDS. For example:

```
//DRVADD EXEC PGM=SLUADMIN
/*
//STEPLIB DD DSN=SYS1.SEALINK
/*
//SLSCNTL DD DISP=SHR,DSN=SYS1.DBASEPRM
//SLSCNTL2 DD DISP=SHR,DSN=SYS1.DBASESEC
//SLSPRINT DD SYSOUT=*
//SLSIN DD *
    SET SLIDRIVS(A00,A01,A02,A03,A04,A05,A06,A07,A08,A09,+
        A10,A11,A12,A13,A14,A15),FORLSMID(00:00),FORPANEL(01)
/*
```

In this example, you are adding two transports to fill the 16-drive panel for LSM 00 (rail 01). A14 and A15 represent the new drives being added; A00-A13 are the existing operational drives.

3. For each ACS affected, use dynamic reconfiguration to bring the new transports online:

```
MODIFY CONFIG UPD ACS(xx) NOCELL
```

Note: If there is more than one ACS to update, wait for the first update to complete before issuing the command to update the next ACS.

The changes cascade to all other hosts.

4. Back up the CDS.

Removing Transports

This section provides a procedure to remove transports from a Streamline Library and HSC configuration. In this case, the Streamline Library panel already contains operational transports, and the HSC remains active during this time, so NCO is used to remove the transports.

To dynamically remove transports from a Streamline Library:

1. Power off the transports and deinstall them.
2. Use the SET SLIDRIVS utility to remove the MVS drive address and drive location to the CDS. For example:

```
//DRVREM EXEC PGM=SLUADMIN
//*
//STEPLIB DD DSN=SYS1.SEALINK
//*
//SLSCNTL DD DISP=SHR,DSN=SYS1.DBASEPRM
//SLSCNTL2 DD DISP=SHR,DSN=SYS1.DBASESEC
//SLSPRINT DD SYSOUT=*
//SLSIN DD *
      SET SLIDRIVS(B00,B01,B02,B03,B04,B05,B06,B07,B08,B09,+
                  B10,B11,B12,B13,,),FORLSMID(00:01),FORPANEL(01)
/*
```

In this example, you are removing two drives and specifying a primary and secondary CDS. B14 and B15 represent the drives being removed; B00-B13 represent the remaining operational drives.

3. For each ACS affected, use dynamic reconfiguration to remove the transports:

```
MODIFY CONFIG UPD ACS(XX) NOCELL
```

Note: If there is more than one ACS to update, wait for the first update to complete before issuing the command to update the next ACS.

The changes cascade to all other hosts.

4. Back up the CDS.

Replacing Transports

The following procedure tells how to replace transports on a Streamline Library panel, while retaining the same unit addresses. The HSC remains active during this time, so NCO is used to replace the transports.

To dynamically replace transports in a Streamline Library:

1. Vary the RTDs that are to be replaced offline, then deinstall them.
2. Install and power on the new transports.
3. Enter MODify CONFIG to replace the transports.
4. Vary the new transports online.
5. Back up the CDS.

Changing Transport Unit Addresses

The following procedure tells how to change the unit addresses for transports on a Streamline Library panel that contains operational drives. The HSC remains active during this time, so NCO is used to change the transport unit addresses.

To dynamically change transport addresses in a Streamline Library:

1. Use the SET SLIDRIVS utility to change the MVS unit address and drive location in the CDS. For example:

```
//UNITADDR EXEC PGM=SLUADMIN
//*
//STEPLIB DD DSN=SYS1.SEALINK
//*
//SLSCNTL DD DISP=SHR,DSN=SYS1.DBASEPRM
//SLSCNTL2 DD DISP=SHR,DSN=SYS1.DBASESEC
//SLSPRINT DD SYSOUT=*
//SYSOUT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//SLSIN DD *
    SET SLIDRIVS(F00,F01,F02,F03,F04,F05,F06,F07,F08,F09,+
                F10,F11,F12,F13,F14,F15),FORLSMID(00:02),FORPANEL(01)
/*
```

In this example, you are changing unit addresses on all transports on a 16-drive panel and specifying a primary and secondary CDS.

2. For each ACS affected, use dynamic reconfiguration to activate the new transport unit addresses:

```
MODIFY CONFIG UPD ACS(XX) NOCELL
```

Note: If there is more than one ACS to update, wait for the first update to complete before issuing the command to update the next ACS.

The changes cascade to all other hosts.

3. Back up the CDS.

Adding SL8500 or SL3000 ACSs

Dynamic hardware reconfiguration for the SL3000 and SL8500 libraries also allows you to dynamically add up to 9 ACSs.

To dynamically add SL8500 or SL3000 ACSs:

1. Run a backup of the CDS.
2. Issue the HSC MODIFY CONFIG ADD ACSType command.

```
MODIFY CONFIG ADD ACSTYPE(SL3000)
```

or

```
MODIFY CONFIG ADD ACSTYPE(SL8500)
```

3. Add a new LMUPATH statement for the ACS.

```
LMUPATH ACS(acs_id) LMUADDR(ip_address)
```

4. Issue the LMUPDEF command on all hosts to activate the new LMUPATH statement.

```
LMUPDEF DATASET(dataset_name)
```

5. Run the SET SLIDRIVS utility or issue the SET SLIDRIVS command to add drives to the new ACS.

```
SET SLIDRIVS(drv1,drv2,drv3,...) FORLSMID (lsm_id), +  
FORPANEL(panel_id)
```

6. Vary the new ACS online to all hosts:

```
VARY ACS aa ONLINE
```

7. Vary the LSMs online.

```
MODIFY LSM lsm_id ONLINE
```

Deleting an SL8500 or SL3000 ACS

Dynamic hardware reconfiguration for the SL3000 and SL8500 libraries also allows you to dynamically delete an entire ACS.

To dynamically delete an SL8500 or SL3000 ACS:

1. Run a backup of the CDS.
2. Vary the LSMs offline.

```
MODIFY LSM lsm_id OFFLINE
```

3. Vary the ACS offline to all hosts.
4. Issue the HSC MODify CONFIG DELete ACSid(*acs-id*) command.
5. Remove the LMUPATH statement defined for the deleted ACS.
6. Issue the LMUPDEF command on all hosts to remove the LMUPATH statement.

Adding and Removing SL8500 or SL3000 Expansion Modules

Adding and removing expansion modules is a hardware-disruptive process. However, HSC can dynamically recognize changes and continue to support other libraries without the need to recycle the HSC.

The following procedure describes how to add or remove expansion modules. The HSC remains active during this time, so NCO is used to make the changes.

To dynamically add or remove SL8500 or SL3000 expansion modules:

1. Enter the MODify LSM OFFline command to modify all Streamline Libraries offline to all HSC hosts in the affected library.

All drives and cartridges in the Streamline Libraries are now unavailable.

2. Physically add or remove the expansion modules.
3. Enter the MODify LSM ONline command to modify the Streamline Libraries online to the HSC.

During the modify online process, HSC automatically discovers the panel configuration changes. A WTOR is issued to the operator describing the change and requesting confirmation to make the changes permanent to the CDS.

- Replying "NO" aborts the configuration change. The CDS is not updated and the old configuration is preserved. This means you could attempt to access a cell that's been removed and any new cell will not be recognized by the HSC.
 - Replying "YES" causes the CDS to be updated with the new configuration and also propagates the changes to all other active hosts in the complex. The HSC updates its in-memory control blocks.
4. Back up the CDS.
 5. Use dynamic reconfiguration to add the expansion modules:

```
MODIFY CONFIG
```

For SL8500s, continue with Step 6. There are no special issues with the SL3000. Run the normal HSC audit, then back up the CDS.

6. Determine the highest panel number in the expanded SL8500.

In an SL8500 library without expansion modules, the highest panel number is 10. Each additional expansion module adds eight panels, so for an SL8500 with one expansion module, the highest panel number is 18; for two expansion modules the highest panel number is 26; and so forth. The three highest panel numbers are in the CEM, so in an SL8500 containing one expansion module, panels 16, 17, and 18 are located in the CEM.
7. Audit the three panels containing the highest panel numbers in each LSM (rail) in the expanded SL8500.
8. Audit any panel that includes cartridges that may have been moved during installation of the Storage Expansion Module.
9. Audit each LSM in the expanded SL8500.
10. Back up the CDS again.

Changing SL8500 CAP Types - HSC Down

The following procedure describes how to change CAP types from SL8500 standard rotational CAPs to a Bulk CAP while the HSC is down.

1. Replace the CAPs in the library and verify they are operational.
2. Start HSC and wait for initialization to complete. See "[Starting and Stopping HSC.](#)"
During startup, reconfiguration starts. Wait for reconfiguration to complete.
3. If the LSMs affected are offline:
 - a. Enter `MODify LSM (lsm-range) ONline.`
where *lsm-range* is a range of LSMs where the CAP changes are being made.
 - b. Enter `Display Cap` to verify the CAP changes to the library.
4. If the LSMs affected are online, enter `Display Cap` to verify the CAP changes to the library.
5. Start HSC activity.

Changing SL8500 CAP Types - HSC Up

The following procedure describes how to change CAP types from SL8500 standard rotational CAPs to a Bulk CAP while the HSC is up.

1. Enter `MODify LSM (lsm-range) OFFline`
where *lsm-range* is the range of LSMs where the CAP changes are being made.
2. Enter `Vary ACS(acs-id) OFFline`
where *acs-id* is the ACS where the CAP changes are to be made.
3. Replace the CAPs in the library and verify they are operational.
4. Stop all HSC activity.
5. Enter `Vary ACS(acs-id) ONline`
where *acs-id* is the ACS where the CAP changes are to be made.
During the vary ACS process, reconfiguration starts. Wait for reconfiguration to complete.
6. Enter `MODify LSM (lsm-range) ONline`
where *lsm-range* is the range of LSMs where the CAP changes are being made.
7. Enter `Display Cap` to verify the CAP changes to the library.
8. Start HSC activity.

Using Audit Instead of MERGEcds for Streamline Libraries

If your environment contains only Streamline Libraries and no virtual or vaulted volume data, you can elect to use the SLUADMIN AUDIT utility to populate the CDS when combining or splitting ACSs instead of using the MERGEcds utility.

The following procedure explains how to use the AUDIt utility to add volumes to a new CDS:

1. Create a new CDS as described in ["Create a New, Empty CDS."](#) Place all LSMs in a single ACS.
2. Enter data into the new CDS as described in ["Prepare to Run MERGEcds"](#) and ["Run the MERGEcds Utility."](#)
3. Ensure that the merge was successful as described in ["Ensure the Merge was Successful."](#)
4. Start HSC with the new CDS on only one host.
Keep the drives offline to MVS, since there is no reason to support automated mounts until AUDIt has added the volumes to the new CDS.
5. Run the AUDIt utility to audit the newly configured LSMs. The order in which the LSMs are audited does not matter, and the entire ACS or all LSMs can be audited at the same time. AUDIt adds the volumes to the new CDS.
6. Run the BACKUp utility to back up the CDS after the audits.
7. Restart HSC on all hosts. Vary the drives online to MVS.
8. Run the Database Decompile (LIBGEN) utility to generate LIBGEN macro statements from a modified CDS.
9. You should then perform a scratch synchronization using either the LCM scratch synchronization function or the SLUCONDDB utility.

Note: If your environment includes virtual or vaulted volume data, you must use the MERGEcds utility to put this data into the new CDS. You can, however, elect to bypass merging the real volumes by using the NOREAL control statement in the SLSMERGE input, and use the SLUADMIN AUDIT utility to populate the real volume data. Using this procedure will result in loss of volume metadata.

Adding a New VTSS

To add VTSSs:

1. Install the new VTSS.
2. Run the DECOMP utility to decompile your CONFIG deck. For more information, refer to the *ELS Command, Control Statement, and Utility Reference*.
3. Update the CONFIG deck with the appropriate VTSS and VTD statements. For more information, refer to the *ELS Command, Control Statement, and Utility Reference* and "[Configuring VTCS](#)."
4. Run the CONFIG utility.
5. If necessary, add volumes.

VTVs and MVCs can be dynamically added or deleted from your configuration by using the SET VOLPARM utility.

HSC SET VAULTVOL and SET VAULT Utilities

The HSC SET VAULTVOL utility can be used to reserve a section of an existing HSC CDS for vaulted volumes. You must ensure that your existing CDS is large enough to hold the projected vault data, using the following formula: The HSC SET VAULT utility can be used to dynamically add, rename, or delete vaults.

If your CDS contains VAULT information, you must specify whether or not this data is to be merged into the new CDS during a CDS merge process.

HSC SET CDKLOCK Utility

The HSC SET CDKLOCK utility can be used to reserve a section of an existing HSC CDS for Open Systems resource lock information. This command can also be used to manage the status of the subfile and change the default expiration of active locks.

Note: The CDS merge process does not copy CDKLOCK information from the old CDS to the new CDS. Resource lock information is transient and may not be valid during the merge process.

Partitioning Streamline Libraries

There are many businesses uses for partitioning your Streamline libraries, for example:

- Enabling service centers to keep a client from accessing another client's tapes,
- Keeping various departments within the same company separate, and
- Dividing a library between production and test environments.

The SL3000 and SL8500 Libraries can be configured with one to eight partitions as ACSs with cell, drive, and CAP assignments. Both Libraries use the StreamLine Library Console (SLC) to perform partition and resource assignment functions.

The following procedures describe how to configure HSC partitioning on the SL3000 or SL8500 library.

- starting from scratch – LIBGEN, SLICREAT
- converting a non-partitioned ACS to a partitioned ACS
- converting a partitioned ACS to a non-partitioned ACS
- adding a library to the HSC Complex
- removing a library from the HSC Complex
- adding a partition to the library
- removing a partition from the library
- moving a partition from one host group to another host group
- adding resources to the library
- removing resources from the library

Partitioning Requirements and Prerequisites

- Read the *SL8500 Best Practices Guide* or *SL3000 User's Guide*.
- Call Oracle Software Support to schedule a CSE to configure the Streamline Library. **This is required.**
- **For SL3000s:** Make sure you have installed firmware Version 2.00 or higher and SLConsole version FRS 4.00 or higher.
- **For SL8500 enhanced partitioning support** (dynamically adding and removing library resources of partitions and so forth), ensure that you have the following prerequisites:
 - SL8500 firmware Compatibility 23 level –Release 7.0x.
 - SLC release level 5.55 or higher

Note: Customers **MUST** be current on maintenance.

- HSC 6.2 (MVS) – L1H16EN and subsequent PTFs
- HSC 6.2 (VM) – L1H16EM and subsequent PTFs
- ELS 7.0 (MVS) – L1H16EO and subsequent PTFs
- ELS 7.1 (MVS) – L1H16EP and subsequent PTFs
- LibraryStation 6.2 - L1S1076
- MVS-CSC 6.2 - 1C109C

Restrictions

- Hosts prior to HSC 6.1, or hosts without the partitioning PTF, cannot bring the partitioned ACS online.

- If the Library being partitioned is not part of the CDS (either as a real or future library), it cannot be dynamically added to the CDS. A LIBGEN/SLICREAT/MergeCDS process must be performed. For more information, see "[Creating the HSC CDS.](#)"
- Volumes remaining in a partition being removed from one host group will be deleted from that CDS when the ACS is deleted using the HSC command Modify Config Delete *acs-id*.
- Volumes located in a partition that is added to a host group will be available to that CDS after an AUDIt or MERGEcds.
- Volume history will only be preserved if MERGEcds is run.
- When converting a partitioned library to non-partitioned library, a MERGEcds is possible to retain volume history from previously defined partitions.
- If a volume is mounted on a drive in a partition that is removed before the drive goes OFFLINE, the volume will be deleted from the CDS, and any future MOUNT activity to that drive will end in an error.
- When moving volumes from one partition to another, neither the HSC CDS nor the library support duplicate volumes.
- Prior to backing up the CDS described in the procedures below, NCS and VTCS systems must be in an idle state. This means draining all batch tape initiators, varying all automated MVS tape devices offline (real and virtual) and varying all RTDs defined in the VTCS configuration offline to VTCS.
- When a VTCS CONFIG RESET is required, all host systems must be down.
- If an HSC group owns multiple partitions within a Streamline library, anything done to that partition is considered disruptive and will affect all partitions connected to an HSC group. For any changes that involve adding or removing partitions or adding or removing partition resources, all activity (i.e., drive(s), cap(s), etc.) must be stopped for all ACSs connected to the library. All ACSs connected to the library must be varied offline. Perform the specific activity and then vary the ACSs online. Activity can now be started for all ACSs.

CAP Considerations

Partitioning affects CAP usage. Consider two situations that can be a common occurrence:

- AUTO CAP – In partitioning mode, AUTO CAP is disabled. If you are using AUTO CAP prior to partitioning the library, change it to manual. If it is not changed, HSC changes it.
- Release CAP – A CAP can only be released from the owning host.
- Overriding a CAP Reservation – Refer to the *SL8500 Best Practices Guide* or *SL3000 User's Guide*.

LibraryStation Considerations

LibraryStation must be recycled to pick up the new drive configuration for the following conditions:

- adding ACSs
- deleting ACSs
- adding LSMs

- deleting LSMs
- adding drives
- deleting drives.

Caution: LibraryStation does not support Near Continuous Operation (NCO). If LibraryStation is running during NCO, results will be unpredictable.

LibraryStation should not be running when NCO is started or is anticipated to start, for example, when you are varying a newly reconfigured library online. This includes changes to drives that are configured to LibraryStation.

Definitions

The following terms are defined here:

- Host Group
- SL8500 Enhanced Partition
- SL3000 Partition

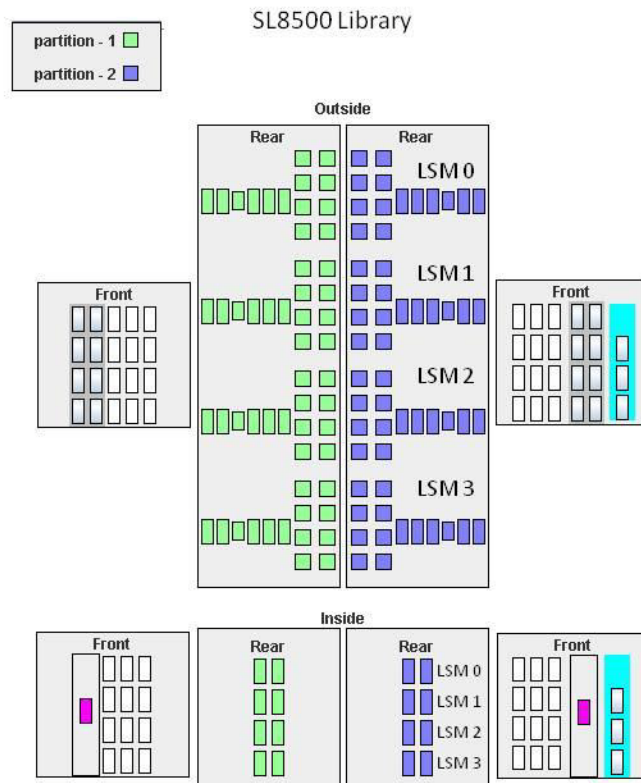
Host Group

A set of hosts connected to an ACS and sharing one CDS. For HSC, there can be up to 16 hosts in a host group.

SL8500 Enhanced Partition

For an SL8500 library, a partition is a segment of one SL8500 unit consisting of cells with a minimum of 1 to n magazines (13 or 14 cells to a magazine), maximum of 2 shared CAPs, and 1 to n drives for a maximum of 8 partitions. A partition is an ACS. A partition has exclusive control of cells and drives. Installed CAPs are shared amongst all defined partitions. [Figure 5-7](#) shows an SL8500 divided into 2 partitions.

Figure 5–7 SL8500 Partitioning Example



SL3000 Partition

For an SL3000 library, a partition is a segment of one SL3000 unit consisting of cells with a minimum of 1 to n in increments of 100, CAPs, and drives for a maximum of 8 partitions. A partition is an ACS. A partition has exclusive control of cells and drives. Installed CAPs may be assigned to one partition (exclusive), assigned to multiple partitions (shared), or unassigned.

Figure 5–8 and Figure 5–9 illustrate an SL3000 divided into 4 partitions:

- Partition 1 assigned to ACSLS:
 - 200 cells
 - 1st and 2nd CAPs shared
 - 24 drives
- Partitions 2 and 3 assigned to HSC Host Group A:
 - 200 cells each
 - 1st and 2nd CAPs shared for each
 - 6 drives each
- Partition 4 assigned to HSC Host Group B:
 - 200 cells

- 1st and 2nd CAPs shared, 3rd and 4th CAPs exclusively assigned
- 12 drives

Figure 5–8 SL3000 Partition and Host Group Example for Cells and CAPs


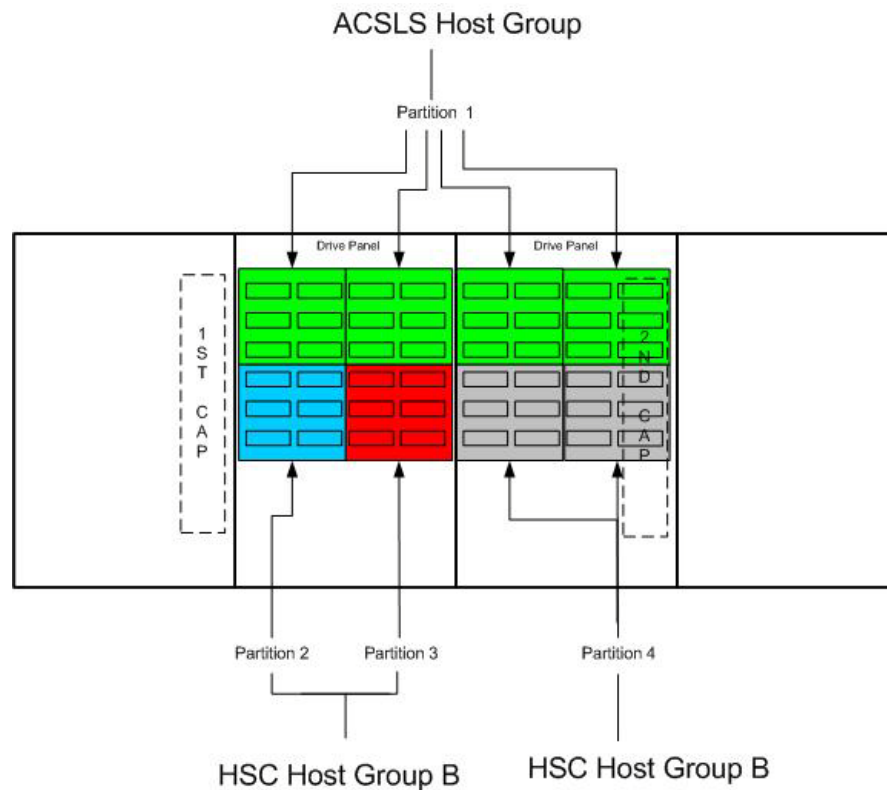


Figure 5–9 SL3000 Partition and Host Group Example for Drives



Partitioning Procedures

The following partitioning procedures are discussed:

- "Starting from Scratch – LIBGEN, SLICREAT"
- "Converting a Non-partitioned Library to a Partitioned Library"
- "Converting a Partitioned Library to a Non-partitioned Library"
- "Adding a Library to the HSC Complex"
- "Removing a Library from the HSC Complex"
- "Adding a Partition to the Library"
- "Removing a Partition from the Library"
- "Moving a Partition from One Host Group to another Host Group"
- "Adding Resources to a Library"
- "Removing Resources from a Library"
- "Varying the ACS Online After a Configuration Change"

Starting from Scratch – LIBGEN, SLICREAT

This procedure creates a partitioned Streamline library from scratch. No LIBGEN changes are required.

Note: It is important that all SL8500 LSMs are online even if they are not defined/allocated to the partition. If the LSMs are not online, CAPs may not be operational.

1. Run an HSC 7.0 or later LIBGEN assemble and the SLICREAT program.

Create a new CDS using the procedures in "[Creating the HSC CDS.](#)" HSC must use HSC 7.0 or later SMP/E libraries. After running the SLICREAT program, a standard CDS is created.

- If virtual tape is being implemented, VTCS CONFIG must be run to define the configuration to VTCS.
- If you plan to add SL3000 or SL8500 ACSs (partitions) in the HSC Complex, it is important to add FUTRACS to the SLILIBRY LIBGEN macro. The first example shows a LIBGEN example for an SL3000 library while the second example shows a LIBGEN example for an SL8500 library.

```

*
LIBGEN  SLIRCVRY  TECHNIQ=SHADOW
*
SLILIBRY  SMF=245,                                X
          ACSLIST=ACSLIST,                          X
          HOSTID=(HST1,HST2),                        X
          DELDISP=SCRTCH,                            X
          MAJNAME=STKASBADD,                         X
          CLNPRFX=CLN,                               X
          COMPRFX=/,                                 X
          FUTRACS=(2),                               X
          SCRLABL=SL                                 X
*
ACSLIST  SLIALIST  ACS00
*
ACS00    SLIACS  LSM=(LSM0000)
*
LSM0000  SLILSM  TYPE=3000
*
*
SLIENDGN ,

*
LIBGEN  SLIRCVRY  TECHNIQ=SHADOW
*
SLILIBRY  SMF=245,                                X
          ACSLIST=ACSLIST,                          X
          HOSTID=(HST1,HST2),                        X
          DELDISP=SCRTCH,                            X
          MAJNAME=STKASBADD,                         X
          CLNPRFX=CLN,                               X
          COMPRFX=/,                                 X
          FUTRACS=(2),                               X
          SCRLABL=SL                                 X
*
ACSLIST  SLIALIST  ACS00
*
ACS00    SLIACS  LSM=(LSM0000,LSM0001,LSM0002,
LSM0003,LSM0004,LSM0005,LSM0006,LSM0007)
*

```

```

LSM0000  SLILSM TYPE=8500
LSM0001  SLILSM TYPE=8500
LSM0002  SLILSM TYPE=8500
LSM0003  SLILSM TYPE=8500
LSM0004  SLILSM TYPE=8500
LSM0005  SLILSM TYPE=8500
LSM0006  SLILSM TYPE=8500
LSM0007  SLILSM TYPE=8500
*
SLIENDGN ,

```

2. Run the ["Converting a Non-partitioned Library to a Partitioned Library"](#) to partition the library.

Note: After the HSC starts for the first time for a newly added SL8500 library, modify all LSMs online even though a partition may not have all the LSMs allocated/defined. For example:

```
MODIFY LSM lsm-id ONLINE
```

The Streamline library is now partitioned and is available for use. Start HSC as usual.

Converting a Non-partitioned Library to a Partitioned Library

This procedure converts an existing non-partitioned Library to a partitioned Library. For SL8500s, the ACS must contain only one SL8500.

After completing the following steps, you will create a partitioned Streamline library.

1. Stop activity for the ACS connected to the library.

In preparation for reconfiguring the partition, stop all CAP and tape activity (both real and virtual) from all active hosts in that library, including RTD tape processing. Failure to do this will result in lost cartridges and errors on dismount.

Vary drives offline to ALL MVS hosts that will no longer have access to the drives. From the MVS console, issue the following command to see if the drives in the affected LSM are offline to MVS:

```
D U,,,xxx-yyy
```

where *xxx-yyy* is the drive ID or UCB range. If the drives are not offline to MVS, vary the drives offline with the MVS command:

```
VARY xxx-yyy OFFLINE
```

Additionally, any RTDs that reside on the affected LSM will need to be varied offline. To do this, issue the VTCS command:

```
V RTD(rtd_id) OFF
```

2. Issue the Display Cap command.

Issue the Display Cap command to verify the CAP is in manual status since this procedure does not go through the NCO process. If the CAP status is AUTOMATIC, issue the CAP Preference command to change the CAP to MANUAL:

```
CAPP prefval cap-id MANUAL
```

3. Vary the affected ACS offline to all hosts. The ACS to be partitioned must be offline to all hosts. Issue the HSC command:

```
VARY ACS acs-id OFFLINE
```

4. Back up the CDS.

Before making changes to the CDS, it is best practice to create a backup copy for recovery purposes. This backup copy can be restored, if problems are encountered, to recover to a known CDS state using documented HSC recovery procedures.

Use the HSC BACKUp utility to create the backup CDS copy.

Note: If VTCS is involved, **all tape activity (real and virtual) must be stopped.** All VTDs must be varied offline to MVS. All RTDs defined in the VTCS configuration must be varied offline to VTCS. It is highly recommended that all initiators be drained that feed both real and VTD processing. NCS and VTCS must be in an idle state prior to creating the backup copy of the CDS to avoid possibility of data loss.

5. Change the Streamline library configuration file to partition the ACS.

The internal configuration file in the Streamline library must be changed must be changed to partition the LSM, assign cell, drive, and CAP resources to that partition. This change is made through the SLConsole.

6. Run the SET SLIDRIVS utility on the host group (optional).

If not all the drives need to be configured to the partitioned ACS, modify drives to the LSM of the ACS by running the SLUADMIN utility for SET SLIDRIVS.

```
SET SLIDRIVS(drv1,drv2,drv3,,,,,,,,,,,,,) + FORLSMID(lsm-id),FORPANEL(12)
```

7. Modify the HSC LMUPATH parameter in PARMLIB.

Add the PARTID(00*n*) keyword to the LMUPATH HSC parameter in PARMLIB. The ID(*n*) is obtained from the SLConsole. For example:

```
LMUPATH ACS(acs-id) LMUADDR(ip-address) PARTID(00n)
```

8. Issue the HSC LMUPDEF command on all hosts.

To activate the modified LMUPATH statement, issue the HSC command:

```
LMUPDEF DATASET(dataset.name)
```

9. Vary the affected ACS online to all hosts. Issue the HSC command:

```
Vary ACS acs-id Online
```

HSC retains the online/offline status of each pre-existing LSM. If the LSMs do not modify online when the ACS is varied online, issue the HSC command:

```
MODIFY LSM lsm-id ONLINE
```

No additional mounts occur to the ACS until the ACS is varied online.

Note: Virtual tape environments are required to vary any RTDs online to VTCS. Additionally, any VTDs varied offline need to be varied online to MVS.

10. Vary drives ONLINE to MVS (if necessary). From the MVS console, issue:

```
D U,,,xxx-yyy
```

where *xxx-yyy* is the drive ID range (both real and virtual), to see if the new drives are online to MVS. If they are not, bring the real drives online with the MVS command:

```
VARY xxx-yyy ONLINE
```

11. Vary RTDs ONLINE to VTCS (if necessary). From the MVS console, issue:

```
V RTD(rtid_id) ON
```

for the RTDs that reside on the LSMs being allocated.

12. Re-establish connections to all defined libraries (if necessary).

Re-establish connections to all defined libraries to acquire drive configuration information from these libraries. To re-drive any outstanding mounts that may have been lost due to hardware or communication errors, issue the SMC RESYNCHRONIZE command.

The ACS is now partitioned and available for use with all four LSMs in a single partition. If you are running multiple HSC hosts, the new configuration is propagated automatically to all hosts.

Converting a Partitioned Library to a Non-partitioned Library

This procedure converts an existing partitioned Library to a non-partitioned Library. For SL8500s, the ACS must contain only one SL8500.

1. Stop activity for the ACS connected to the library.

Caution: In preparation for reconfiguring the partition, stop all CAP and tape activity (both real and virtual) from all active hosts in that library, including RTD tape processing. **Failure to do this will result in lost cartridges and errors on dismount.**

Vary drives offline to ALL MVS hosts that will no longer have access to the drives. From the MVS console, issue the following command to see if the drives in the affected LSM are offline to MVS:

```
D U,,,xxx-yyy
```

where *xxx-yyy* is the drive ID or UCB range. If the drives are not offline to MVS, vary the drives offline with the MVS command:

```
VARY xxx-yyy OFFLINE
```


Additionally, any RTDs that reside on the affected LSM will need to be varied offline. To do this, issue the VTCS command:

```
V RTD(rtid_id) OFF
```

Caution: Failure to perform this step may result in the devices being removed becoming allocated. This may cause **severe operational consequences** including loss of access for any cartridge mounted at the time of the non-partitioning.

If VTCS is involved, any RTDs moving out of the CDS of the host groups affected will require a VTCS CONFIG RESET to be run to remove the specific RTDs from the VTCS configuration.

2. Run the Volume Report utility.

Run a volume report on the affected ACS(s) and keep the output. This data can be used later to verify where volumes were in the ACS(s) before it was removed.

Systems with VTCS should run an MVC report for cross reference purposes.

3. Back up the CDS.

Before making changes to the CDS, it is best practice to create a backup copy for recovery purposes. This backup copy can be restored, if problems are encountered, to recover to a known CDS state using documented HSC recovery procedures.

Use the HSC BACKUP utility to create the backup CDS copy.

Note: If VTCS is involved, all tape activity (real and virtual) must be stopped. All VTDs must be varied offline to MVS. All RTDs defined in the VTCS configuration must be varied offline to VTCS. It is highly recommended that all initiators be drained that feed both real and VTD processing. **NCS and VTCS must be in an idle state prior to creating the backup copy of the CDS to avoid the possibility of data loss.**

4. Vary all ACSs connected to the library OFFLINE to all host groups.

The ACS(s) to be un-partitioned must be offline to all hosts. Issue the HSC command:

```
VARY ACS acs-id OFFLINE
```

5. Change the Streamline library configuration file to unpartition the ACS.

The internal configuration file in must be changed to delete all partitions. This change is made through the SLConsole.

6. Issue the HSC F CONFIG DELeTe ACS command.

To delete an ACS from an HSC host group, issue the HSC command as in the examples below:

```
F CONFIG DELETE ACS(03)
F CONFIG DELETE ACS(04)
F CONFIG DELETE ACS(05)
```

Note:

- If running multiple HSC hosts, the new configuration will be propagated automatically to all hosts.
- Deleting an ACS deletes its volumes from the CDS.
- If this is the last ACS in the HSC host group, the following steps may apply. At this time, the HSC can be stopped.
 - Run a volume report on the affected ACS and keep the output. This data can be used later to verify where volumes were in the ACS before it was removed.
 - Eject all cartridges from the affected ACS.

7. Run the SET SLIDRIVS utility on the remaining host group.

Add drives to the new LSM of the ACS by running the SLUADMIN utility or command for SET SLIDRIVS.

```
SET SLIDRIVS(drv1,drv2,drv3,,,,,,,,,,,,,) +
FORLSMID(lsm-id),FORPANEL(12)
```

8. Modify LMUPDEF in PARMLIB for all host groups.

Remove all the LMUPATH statements referencing the ACSs previously defined to all LMUPDEF HSC parameter files of all affected host groups. If a particular host group will have ownership of the library after it becomes un-partitioned, add or change the LMUPATH statement using the following example:

```
LMUPATH ACS(03) LMUADDR(ip-address)
```

9. Issue the HSC LMUPDEF command on all host groups. To activate the modified HSC LMUPATH statement, issue the HSC command:

```
LMUPDEF DATASET(dataset.name)
```

10. Vary the ACS ONLINE to the remaining host group. Issue the HSC command for the host group taking ownership of library:

```
Vary ACS acs-id Online
```

The HSC retains the online/offline status of each pre-existing LSM. If the LSM does not modify online when the ACS is varied online, issue the HSC command:

```
MODIFY LSM lsm-id ONLINE
```

No additional mounts will occur to the ACS until the ACS is varied online. All other hosts will complete their configuration process when the ACS is varied online.

Virtual tape environments will be required to vary any RTDs online to VTCS.

Note: When an ACS is varied online, the panels are discovered to complete the configuration process. During this time, if any cartridges are discovered, an SLS4416E message will be displayed. It is important to follow the user response section of this message before proceeding to the next step. This section will instruct the user to perform the following step in the order listed:

- Run the HSC SLUADMIN Audit utility with APPLY(NO) for the panels indicated to discover and resolve issues.
 - Run the HSC SLUADMIN Volume Report utility if necessary to gather volume information.
 - Run the HSC SLUADMIN Audit utility with APPLY(YES) to apply updates to the CDS.
 - Run the HSC SLUADMIN SET Utility specifying FREEZE OFF for any panels in the frozen state.
-

Refer to "[Varying the ACS Online After a Configuration Change](#)" for further details of this process.

11. Vary drives ONLINE to MVS (if necessary). From the MVS console, issue the following command to see if the new drives are online to MVS:

```
D U, , xxx-yyy
```

where *xxx-yyy* is the drive ID range (both real and virtual). If they are not, bring the real drives online with the MVS command:

```
VARY xxx-yyy ONLINE
```

12. Vary RTDs ONLINE to VTCS (if necessary). From the MVS console, issue

```
V RTD(rtd_id) ON
```

for the RTDs that reside on the LSMs being allocated.

13. Run the HSC Scratch Update utility (if necessary).

This step can be bypassed if no volumes have been added, deleted, or moved. It must be run if new LSMs have been added to the configuration.

14. Re-establish connections to all defined libraries (if necessary) to acquire drive configuration information from these libraries. To re-drive any outstanding mounts that may have been lost due to hardware or communication errors, issue the SMC RESYNCHRONIZE command.

The ACS is now unpartitioned and available for use. If you are running multiple HSC hosts, the new configuration is propagated automatically to all hosts.

Adding a Library to the HSC Complex

This procedure adds a SL3000 or SL8500 Library for either an existing or new configuration. After completion of the following steps, a SL3000 or SL8500 library is ready for use.

To add a library to the HSC Complex:

1. Back up the CDS (if necessary).

Issue the following command to modify the LSM online:

```
MODIFY LSM lsm-id ONLINE
```

All other hosts will complete their configuration process when the ACS is varied online.

8. Audit LSM of ACS (optional).

From Step 7, when an SLS4416E message is displayed, some panels are populated with cartridges that are not cataloged to HSC. These panels must be audited for HSC to be aware of these cartridges. To audit the panel, run the HSC SLUADMIN AUDIT utility as follows:

```
AUDIT ACS(xx) LSM(yy) PANEL(zz)
```

9. Unfreeze panels/LSM (optional).

If the panels were audited in Step 8, the panels can now be unfrozen. To unfreeze the panels/LSM, run the HSC SLUADMIN utility for SET FREEZE as follows:

```
SET FREEZE (OFF) FORLSMID(lsm-id) FORPANEL(panel-id)
```

10. Run the HSC SCRATCH Update utility (if necessary).

This step can be bypassed if no volumes have been added, deleted, or moved. It must be run if new LSMs have been added to the configuration.

11. Vary drives ONLINE to MVS (if necessary). From the MVS console, issue:

```
D U,,,xxx-yyy
```

where *xxx-yyy* is the drive id range (both real and virtual), to see if the new drives are online to MVS. If they are not, bring the real drives online with the MVS command:

```
VARY xxx-yyy ONLINE
```

12. Vary RTDs ONLINE to VTCS (if necessary). From the MVS console, issue:

```
V RTD(rtd_id) ON
```

for the RTDs that reside on the LSMs being allocated.

Note:

- Virtual tape environments will be required to vary any RTDs online to VTCS. Additionally, any virtual tape drives varied offline will need to be varied online to MVS.
 - Prior to performing this step, ensure that SET SLIDRIVS (Step 6) is performed so that the CDS drive configuration matches EXACTLY to the new drive configuration on the SL3000 Library.
-
-

13. Re-synchronize SMC. To add drives for SMC, issue the SMC RESYNCHRONIZE command.

The Library is now added to the HSC complex and is available for use to all HSC hosts. If running multiple HSC hosts, the new configuration will be propagated automatically to all hosts.

Removing a Library from the HSC Complex

This procedure removes a Library from the HSC Complex.

Note: If this is the last ACS in the HSC host group, only Steps 1 through 6 may apply. At this time, HSC can be stopped.

To remove a library from the HSC Complex:

1. Stop activity for affected ACS connected to HSC.

In preparation for removing a Library from HSC, stop all affected ACS CAP and tape activity (both real and virtual) from all active hosts, including RTD tape processing, in that ACS. Failure to do so will result in lost cartridges and errors on dismount.

Vary drives (both real and virtual) offline to all MVS hosts that will no longer have access to the drives. From the MVS console, issue:

```
D U,,,xxx-yyy
```

where *xxx-yyy* is the drive ID or UCB range, to see if the drives in the affected ACS are offline to MVS. If not, vary the drives offline with the MVS command:

```
VARY xxx-yyy OFFLINE
```

Additionally, any RTDs that reside on the affected ACS will need to be varied offline. To do this, issue the VTCS command:

```
V RTD(rtd_id) OFF
```

Caution: Failure to perform this step may result in the devices being removed becoming allocated. This may cause severe operational consequences including loss of access for any cartridge mounted at the time of the partitioning.

If VTCS is involved any RTDs moving out of the CDS will require a VTCS CONFIG RESET to be run to remove the specific RTDs from the VTCS configuration.

2. Run the SET Freeze utility to freeze affected LSM.

To stop cartridges from migrating to the LSM of the partition being removed, run the HSC SLUADMIN SET FREEZE utility:

```
SET FREEZE(ON) FORLSMID(lsm-id)
```

Note:

- When an LSM is frozen, no new cell allocations will occur, however, any cartridge from the LSM that is currently mounted on a drive will return to its home cell.
 - Since there is one LSM to an SL3000 Library, freezing the LSM is freezing the ACS.
-
-

3. Run the Volume Report utility.

Run a volume report on the affected ACS and keep the output. This data can be used later to verify where volumes were in the ACS before it was removed.

Note: Systems with VTCS should run an MVC report for cross reference purposes.

4. Eject cartridges out of the affected ACS (optional). Eject all cartridges from the affected ACS.**5.** Vary affected ACS connected to HSC OFFLINE to all hosts.

The affected ACS connected to HSC must be offline to all hosts. Issue the HSC command:

```
VARY ACS acs-id OFFLINE
```

6. Back up the CDS.

Before making changes to the CDS, it is best practice to create a backup copy for recovery purposes. This backup copy can be restored if problems are encountered to recover to a known CDS state using documented HSC recovery procedures.

Use the HSC BACKUP utility to create the backup CDS copy.

Note: If VTCS is involved, ALL tape activity (real and virtual) must be stopped. All virtual tape drives must be varied offline to MVS. All RTDs defined in the VTCS configuration must be varied offline to VTCS. It is highly recommended that all initiators be drained that feed both real and virtual tape drive processing. NCS and VTCS must be in an idle state prior to creating the backup copy of the CDS to avoid possibility of data loss.

7. Issue the HSC F CONFIG DELeTe ACS command.

To delete a partition(s) in an HSC Host Group, issue the HSC command as the examples below show:

```
F CONFIG DELETE ACS(03)
F CONFIG DELETE ACS(04)
F CONFIG DELETE ACS(05)
```

Note:

- If running multiple HSC hosts, the new configuration will be propagated automatically to all hosts.
 - Deleting an ACS deletes its volumes from the CDS. Remove the HSC LMUPATH parameter in PARMLIB.
-

8. Remove the LMUPATH statements referencing the ACS previously defined to the LMUPDEF HSC parameter file of the affected host group.

9. Issue the HSC LMUPDEF command on all hosts.

To refresh the modified HSC LMUPDEF HSC parameter file within a host group, issue the HSC command to all hosts:

```
LMUPDEF DATASET(dataset.name)
```

10. Re-synchronize SMC. Issue the SMC RESYNCHRONIZE command.

The Library is removed from the HSC Complex.

Adding a Partition to the Library

This procedure adds a partition to the library for either an existing or new configuration.

After completing the following steps, a partitioned SL3000 or SL8500 with up to eight partitions may be created. At a later time, remove any partitions from this library using "[Removing a Partition from the Library](#)."

If you are adding more partitions to an ACS than what currently exist, the FUTRACS parameter in the SLILIBRY macro must be coded in the LIBGEN with enough ACSs to accommodate the total number of partitions being accessed within a SL3000 or SL8500 library.

1. Stop activity for all ACSs connected to the library.

In preparation for adding a partition to the library, stop all CAP and tape activity (both real and virtual) from all active hosts, including RTD tape processing, in that ACS. **Failure to do so will result in lost cartridges and errors on dismount.**

Vary drives (both real and virtual) offline to all MVS hosts that will no longer have access to the drives. From the MVS console, issue:

```
D U , , ,xxx-yyy
```

where *xxx-yyy* is the drive ID or UCB range, to see if the drives in the affected ACS are offline to MVS. If not, vary the drives offline with the MVS command:

```
VARY xxx-yyy OFFLINE
```

Additionally, any RTDs that reside on the affected ACS will need to be varied offline. To do this, issue the VTCS command:

```
V RTD(rtid_id) OFF
```

Note:

- This step is necessary if you are adding partitions other than the first time.
 - If VTCS is involved, any RTDs moving out of the CDS will require a VTCS CONFIG RESET to be run to remove the specific RTDs from the VTCS configuration.
-
-

2. Vary all ACSs connected to the library OFFLINE to all hosts. ALL ACSs connected to the library must be offline to all hosts. Issue the HSC command:

```
VARY ACS acs-id OFFLINE
```

3. Back up the CDS (if necessary).

Before making changes to the CDS, it is best practice to create a backup copy for recovery purposes. This backup copy can be restored if problems are encountered to recover to a known CDS state using documented HSC recovery procedures.

Use the HSC BACKUP utility to create the backup CDS copy.

Note:

- This step is necessary if you are adding partitions other than the first time.
 - If VTCS is involved, **ALL tape activity (real and virtual) must be stopped**. All virtual tape drives must be varied offline to MVS. All RTDs defined in the VTCS configuration must be varied offline to VTCS. It is highly recommended that all initiators be drained that feed both real and virtual tape drive processing. NCS and VTCS must be in an idle state prior to creating the backup copy of the CDS to avoid possibility of data loss.
-
-

4. Change the SL3000 or SL8500 configuration to add the partition and assign resources to the partition.

The internal configuration file in the SL3000 or SL8500 must be changed to show the newly added partition(s) to a library. This change is made through the SLConsole.

Note: If scratch volumes exist, update and reload SCRDEF, VOLDEF, and TAPEREQ PARMLIB members accordingly. This must be done prior to executing Step 5.

5. Issue the HSC F CONFIG ADD ACSType command (optional).

If you are adding more partitions to an ACS than currently exist, issue the HSC command as in the examples below:

```
F CONFIG ADD ACST(SL3000)
F CONFIG ADD ACST(SL8500)
```

6. Change or add the HSC LMUPATH statement in PARMLIB.

If an LMUPATH statement was defined for an SL3000 or SL8500 using "[Starting from Scratch – LIBGEN, SLICREAT](#)," change the HSC LMUPATH statement to add the PARTID keyword in LMUPDEF in PARMLIB. Verify the partition ID defined from the SLConsole is the same as the PARTID in the HSC LMUPDEF. For a new partition, add an HSC LMUPATH statement in LMUPDEF in PARMLIB.

For example:

```
LMUPATH ACS(acs-id) LMUADDR(ip-address) PARTID(nnn)
```

7. Issue the HSC LMUPDEF command on all hosts. To activate the modified HSC LMUPATH statement, issue the HSC command:

```
LMUPDEF DATASET(dataset.name)
```

8. Run the SET SLIDRIVS utility (match CDS drive configuration to the library). Add drives to an LSM of an added ACS by running the SLUADMIN utility or command for SET SLIDRIVS.

```
SET SLIDRIVS(drv1,drv2,drv3,,,,,,,,,,,,,) +  
FORLSMID(lsm-id),FORPANEL(12)
```

Note: If VTCS is involved, any RTDs moving into the CDS will require a VTCS CONFIG RESET to be run to add the specific RTDs to the VTCS configuration.

9. Vary all ACSs connected to the library ONLINE to all hosts. Issue the HSC command:

```
Vary ACS acs-id Online
```

The HSC retains the online/offline status of each pre-existing LSM. If the LSM is not modified online when the ACS is varied online, issue the HSC command:

```
MODIFY LSM lsm-id ONLINE
```

No additional mounts will occur to the ACS until the ACS is varied online. All other hosts will complete their configuration process when the ACS is varied online.

10. Audit LSM of ACS (optional).

From Step 9, when an SLS4416E message is displayed, some panels are populated with cartridges that are not cataloged to HSC. These panels must be audited for the HSC to be aware of these cartridges. To audit the panel, run the HSC AUDIT utility as follows:

```
AUDIT ACS(xx) LSM(yy) PANEL(zz)
```

11. Unfreeze panels/LSM (optional).

If the panels were audited in Step 10, the panels can now be unfrozen. To unfreeze the panels/LSM, run the HSC SLUADMIN SET FREEZE utility as follows:

```
SET FREEZE (OFF) FORLSMID(lsm-id) FORPANEL(panel-id)
```

12. Run the HSC Scratch Update utility (if necessary). Run an HSC Scratch Update utility.

This step may be bypassed if no volumes have been added, deleted, or moved. It must be run if new LSMs have been added to the configuration.

13. Vary drives ONLINE to MVS (if necessary).

From the MVS console, issue the following command to see if the new drives are online to MVS:

```
D U,,,xxx-yyy
```

where *xxx-yyy* is the drive ID range (both real and virtual). If they are not, bring the real drives online with the MVS command:

```
VARY xxx-yyy ONLINE
```

14. Vary RTDs ONLINE to VTCS (if necessary). From the MVS console, issue:

```
V RTD(rtid_id) ON
```

for the RTDs that reside on the LSMs being allocated.

Note:

- Virtual tape environments will be required to vary any RTDs online to VTCS. Additionally, any virtual tape drives varied offline will need to be varied online to MVS.
 - Prior to performing this step, ensure that SET SLIDRIVS (Step 8) is performed so that the CDS drive configuration matches EXACTLY to the new drive configuration on the SL3000 library.
-
-

15. Re-synchronize the SMC. To add drives for the SMC, issue the SMC RESYNCHRONIZE command.

The partition is now added to the library and is available for use. If running multiple HSC hosts, the new configuration will be propagated automatically to all hosts.

Removing a Partition from the Library

Note: If this is the last ACS in the HSC host group, only Steps 1 through 6 may apply. At this time, the HSC can be stopped.

This procedure removes a partition from an existing library.

1. Stop activity for **affected ACSs being deleted**.

In preparation for removing a partition from the library, stop all CAP and tape activity (both real and virtual) from all active hosts, including RTD tape processing, in that ACS. **Failure to do so will result in lost cartridges and errors on dismount.**

Vary drives (both real and virtual) offline to all MVS hosts that will no longer have access to the drives. From the MVS console, issue the following command to see if the drives in the affected ACS are offline to MVS:

```
D U,,,xxx-yyy
```

where *xxx-yyy* is the drive ID or UCB range. If not, vary the drives offline with the MVS command:

```
VARY xxx-yyy OFFLINE
```

Additionally, any RTDs that reside on the affected ACS will need to be varied offline. To do this, issue the VTCS command:

```
V RTD(rtd_id) OFF
```

Caution: Failure to perform this step may result in the devices being removed becoming allocated. This may cause severe operational consequences including loss of access for any cartridge mounted at the time of the partitioning.

If VTCS is involved, any RTDs moving out of the CDS will require a VTCS CONFIG RESET to be run to remove the specific RTDs from the VTCS configuration.

2. Run the SET Freeze utility to freeze affected LSM.

To stop cartridges from migrating to the LSM of the partition being removed, run the HSC SLUADMIN SET FREEZE utility:

```
SET FREEZE(ON) FORLSMID(lsm-id)
```

Note:

- When an LSM is frozen, no new cell allocations will occur, however, any cartridge from the LSM that is currently mounted on a drive will return to its home cell.
 - Since there is one LSM in an SL3000 library, freezing the LSM is freezing the ACS.
-
-

3. Run the Volume Report utility.

Run a volume report on the affected ACS and keep the output. This data can be used later to verify where volumes were in the ACS before it was removed.

Systems with VTCS should run an MVC report for cross-reference purposes.

4. Eject all cartridges from the affected ACS (optional).
5. Vary the affected ACS OFFLINE to all hosts. ACSs being deleted must be offline to all hosts. Issue the HSC command:

```
VARY ACS acs-id OFFLINE
```

6. Back up the CDS.

Before making changes to the CDS, it is best practice to create a backup copy for recovery purposes. This backup copy can be restored if problems are encountered to recover to a known CDS state using documented HSC recovery procedures.

Use the HSC BACKUP utility to create the backup CDS copy.

Note: If VTCS is involved, ALL tape activity (real and virtual) must be stopped. All virtual tape drives must be varied offline to MVS. All RTDs defined in the VTCS configuration must be varied offline to VTCS. It is highly recommended that all initiators be drained that feed both real and virtual tape drive processing. HSC and VTCS must be in an idle state prior to creating the backup copy of the CDS to avoid possibility of data loss.

7. Issue the HSC F CONFIG DELeTe ACS command.

To delete a partition(s) in an HSC Host Group, issue the HSC command as the examples below show:

```
F CONFIG DELETE ACS(03)
F CONFIG DELETE ACS(04)
F CONFIG DELETE ACS(05)
```

Note:

- If running multiple HSC hosts, the new configuration will be propagated automatically to all hosts.
 - Deleting an ACS deletes its volumes from the CDS.
 - If the last allocated ACS is deleted, the HSC no longer has connections to any library and can be stopped. Steps 8 through 10 can be ignored.
-
-

8. Remove the HSC LMUPATH parameter in PARMLIB.

Remove the LMUPATH statements referencing the ACS previously defined to the LMUPDEF HSC parameter file of the affected host group.

9. Issue the HSC LMUPDEF command on all hosts.

To refresh the modified HSC LMUPDEF HSC parameter file within a host group, issue the HSC command to all hosts:

```
LMUPDEF DATASET(dataset.name)
```

10. Change the SL3000 configuration to remove the affected partition from the library (optional).

The internal configuration file in the SL3000 may be changed to remove the partition. This change happens through the SLConsole.

11. Vary all ACSs connected to the library ONLINE to all hosts. Issue the HSC command:

```
Vary ACS acs-id Online
```

The HSC retains the online/offline status of each pre-existing LSM. If the LSM is not modified online when the ACS is varied online, issue the HSC command:

```
MODIFY LSM lsm-id ONLINE
```

No additional mounts will occur to the ACS until the ACS is varied online. All other hosts will complete their configuration process when the ACS is varied online.

12. Vary drives ONLINE to MVS (if necessary). From the MVS console, issue the following command to see if the new drives are online to MVS:

```
D U, , , xxx-yyy
```

where *xxx-yyy* is the drive ID range (both real and virtual). If they are not, bring the real drives online with the MVS command:

```
VARY xxx-yyy ONLINE
```

13. Vary RTDs ONLINE to VTCS (if necessary). From the MVS console, issue:

```
V RTD(rtid_id) ON
```

for the RTDs that reside on the LSMs being allocated.

Note: Virtual tape environments will be required to vary any RTDs online to VTCS. Additionally, any virtual tape drives varied offline will need to be varied online to MVS.

14. Re-synchronize the SMC. Issue the SMC RESYNCHRONIZE command.

The partition is removed from the library, and the remaining library resources are available for use.

Moving a Partition from One Host Group to another Host Group

This procedure moves a partition from one host group to another and preserves the history of the partition volumes being allocated.

Note: If this is the last ACS in the "from" HSC host group, only steps 1 through 4 may apply. At this time, the "from" HSC can be stopped. Proceed with Steps 9 through 20.

"From" Host Group:

1. Stop activity for all ACSs connected to the library.

In preparation for moving the partition, you must stop all CAP and tape activity (both real and virtual), from all active hosts, including virtual tape processing, in that ACS of the host group of where the partition is moving from. **Failure to do this will result in lost cartridges and errors on dismount.**

Vary drives offline to all MVS hosts that will no longer have access to the drives. From the MVS console, issue the following command to see if the drives in the affected LSM are offline to MVS.

```
D U, , , xxx-yyy
```

where *xxx-yyy* is the drive ID or UCB range. If they are not, bring them offline with the MVS command:

```
VARY xxx-yyy OFFLINE
```

Additionally, any RTDs that reside on the affected LSM will need to be varied offline. To do this issue the VTCS command:

```
V RTD(rtid_id) OFF
```

Caution: Failure to perform this step to vary these drives offline from MVS will result in these devices becoming allocated while removing the LSM from the partition. This may cause severe operational consequences including loss of access for any cartridge that was mounted at the time of the partitioning.

If VTCS is involved, any RTDs moving out of the CDS of the "from host group" will require a VTCS CONFIG RESET to be run to remove the specific RTDs from the VTCS configuration.

2. Run the Volume Report utility. Run a volume report on the affected ACS and keep the output. This data can be used later to verify where volumes were in the ACS before it was relocated.

Note: Systems with VTCS should run an MVC report for cross reference purposes.

3. Vary the ACS being removed OFFLINE to all hosts. The ACS must be offline to all hosts. Issue the HSC command:

```
VARY ACS acs-id OFFLINE
```

4. Back up the CDS (CDS01). Use the HSC BACKUP utility to create the backup CDS copy as CDS01.
5. Remove the HSC LMUPATH parameter in PARMLIB. Remove the LMUPATH statements referencing the partition previously defined to the LMUPDEF HSC parameter file.
6. Issue the HSC LMUPDEF command on all hosts. To refresh the modified HSC LMUPDEF HSC parameter file within a host group, issue the HSC command to all hosts:

```
LMUPDEF DATASET(dataset.name)
```

7. Issue the HSC F CONFIG DELeTe ACS command.

To delete a partition from a HSC Host Group, issue the HSC command as in the examples below:

```
F CONFIG DELETE ACS(03)
F CONFIG DELETE ACS(04)
F CONFIG DELETE ACS(05)
```

Note:

- If you are running multiple HSC hosts, the new configuration will be propagated automatically to all hosts.
 - Deleting an ACS deletes its volumes from the CDS.
-
-

8. Re-synchronize the SMC. To add drives for the SMC, issue the SMC RESYNCHRONIZE command.

"To" Host Group:

1. Back up the CDS (CDS02). Use the HSC BACKUP utility to create the backup CDS copy as CDS02.
2. Issue the HSC F CONFIG ADD ACSType command.

To add a partition to an HSC Host Group, issue the HSC command as in the examples below:

```
F CONFIG ADD ACST(SL3000)
F CONFIG ADD ACST(SL8500)
```

Note: If running multiple HSC hosts, the new configuration will be propagated automatically to all hosts.

3. Add the HSC LMUPATH statement in PARMLIB. Add an HSC LMUPATH statement in LMUPDEF in PARMLIB.

For example:

```
LMUPATH ACS(acs-id) LMUADDR(ip-address) PARTID(nnn)
```

4. Issue the HSC LMUPDEF command on all hosts. To activate the modified HSC LMUPATH statement, issue the HSC command:

```
LMUPDEF DATASET(dataset.name)
```

5. Set the Service Level to BASE for all hosts.

Set the service level to BASE for the host group to allow a MERGEcds to be started. Issue the HSC command:

```
SRVLEV BASE
```

6. Run the MERGEcds utility.

Input will be CDS01. Output will be an updated CDS02. This example will move the ACS volume history from ACS 00 (CDS01) to ACS 04.

Caution:

- If you are changing the library configuration, for example, changing UPC labels, the cartridges in the cells being removed must be ejected or moved, and the cells being removed must be unassigned to all partitions using the SLConsole prior to physically adding the library resources.
 - When you are using the SLConsole to add ACS resources, make all necessary changes at the same time and issue one APPLY. The HSC performs a reconfiguration every time an SLConsole partition APPLY occurs.
-

1. Stop activity for all ACSs connected to the library.

In preparation for adding resources to the library, stop all CAP and tape activity (both real and virtual) from all active hosts in that library, including RTD tape processing. **Failure to do this will result in lost cartridges and errors on dismount.**

Vary drives offline to ALL MVS hosts that will no longer have access to the drives. From the MVS console, issue the following command to see if the drives in the affected LSM are offline to MVS.

```
D U, , ,xxx-yyy
```

where *xxx-yyy* is the drive ID or UCB range. If not, vary the drives offline with the MVS command:

```
VARY xxx-yyy OFFLINE
```

Additionally, any RTDs that reside on the affected LSM will need to be varied offline. To do this, issue the VTCS command:

```
V RTD(rtid_id) OFF
```

Caution: Failure to perform this step may result in the devices being removed becoming allocated. This may cause severe operational consequences including loss of access for any cartridge mounted at the time of the non-partitioning.

If VTCS is involved, any RTDs moving out of the CDS of the host groups affected will require a VTCS CONFIG RESET to be run to remove the specific RTDs from the VTCS configuration.

2. Back up the CDS containing the ACS to be changed. Use the HSC BACKUP utility to create the backup CDS copy.

Note: If VTCS is involved, all CAP and tape activity (real and virtual) must be stopped. All virtual tape drives must be varied offline to MVS. All RTDs defined in the VTCS configuration must be varied offline to VTCS.

It is highly recommended that all initiators be drained that feed both real and virtual tape drive processing. NCS and VTCS must be in an idle state prior to creating the backup copy of the CDS to avoid possibility of data loss.

3. Vary all ACSs connected to the library OFFLINE to all hosts. ALL ACSs connected to the library must be offline to all hosts. Issue the HSC command:

```
VARY ACS acs-id OFFLINE
```

4. Change the Library configuration to the affected ACS.

The internal configuration file in the Library is changed to reconfigure the ACS, i.e., adding or deleting cell allocations or changing the exclusivity status of a CAP. This change is accomplished through the SLConsole.

- CAPs:
 - If assigning a CAP, use the SLConsole to assign the CAP to only one partition (dedicated) or a set of partitions (shared).
 - Only dedicated SL3000 CAPs can be set to auto. All SL8500 CAPs in a partitioned library are shared.
 - Prior to reassigning a CAP, it must be set to manual.
 - If installing a CAP, verify the cells that occupied the CAP space are unassigned to any partition. Assign the CAP as described in the previous bullet.
- Drives – If installing or assigning a drive, use the SLConsole to assign the drive to a partition.
- Drive Bays – If installing a drive bay, after verifying the cells are unassigned to any partition, use the SLConsole to assign the drive slots to a partition.
- Cells – If assigning cells, use the SLConsole to assign cells to a partition.
- SL3000:
 - Dual Tallbots – If installing a second robot, verify the outermost three columns (front and back) of each outermost Cartridge Expansion Module (CEM) are unassigned to any partition.
 - Outermost CEMs must be transformed into Parking Expansion Modules (PEMs) when installing the second robot.
 - Expansion Modules – If installing a Drive Expansion Module (DEM) no additional setup is needed. If installing a CEM to either end, no additional setup is needed
 - Prior to installing a CEM to a location other than to either end of the library (that is, between a PEM and either a CEM, DEM, or Base Drive Module) the volumes in the PEM must be either moved to available space within the partition or ejected if no space is available. After the volumes


```
F CONFIG UPDATE ACS(acs-id)
```

where *acs-id* is the ID of the affected ACS.

9. Vary drives ONLINE to MVS if adding drives (optional). From the MVS console, issue the following command to see if the new drives are online to MVS.

```
D U,,,xxx-yyy
```

where *xxx-yyy* is the drive ID range (both real and virtual). If they are not, bring the real drives online with the MVS command:

```
VARY xxx-yyy ONLINE
```

10. Vary drives ONLINE to VTCS if adding RTDs (optional). From the MVS console, issue:

```
V RTD(rtd_id) ON
```

for the RTDs that reside on the LSMs being allocated.

11. Re-synchronize the SMC. Issue the SMC RESYNCHRONIZE command.

If you are running multiple HSC hosts, the new configuration will be propagated automatically to all hosts in the affected host group.

Removing Resources from a Library

This procedure allows resources to be removed from an existing Library or partition of a Library defined as an ACS. Resource removal can consist of cells, drives, drive bays, CAPs, expansion modules and SL3000 dual Tallbot.

Note: When you are using the SLConsole to remove ACS resources, make all necessary changes at the same time and issue one APPLY. The HSC performs a reconfiguration every time an SLConsole partition APPLY occurs.

1. Back up the CDS containing the ACS to be changed. Use the HSC BACKUP utility to create the backup CDS copy.

Note: If VTCS is involved, ALL CAP and tape activity (real and virtual) must be stopped. All virtual tape drives must be varied offline to MVS. All RTDs defined in the VTCS configuration must be varied offline to VTCS. It is highly recommended that all initiators be drained that feed both real and virtual tape drive processing. HSC and VTCS must be in an idle state prior to creating the backup copy of the CDS to avoid possibility of data loss.

2. Stop activity for all ACSs connected to the library.

In preparation for removing resources from the library, stop all CAP and tape activity (both real and virtual) from all active hosts in that ACS, including RTD tape processing. **Failure to do this will result in lost cartridges and errors on dismount.**

Vary drives offline to ALL MVS hosts that will no longer have access to the drives. From the MVS console, issue the following command to see if the drives in the affected LSM are offline to MVS.

```
D U,,,xxx-yyy
```

where *xxx-yyy* is the drive ID or UCB range. If not, vary the drives offline with the MVS command:

```
VARY xxx-yyy OFFLINE
```

Additionally, any RTDs that reside on the affected LSM will need to be varied offline. To do this, issue the VTCS command:

```
V RTD(rtd_id) OFF
```

Caution: Failure to perform this step may result in the devices being removed becoming allocated. This may cause severe operational consequences including loss of access for any cartridge mounted at the time of the partitioning.

If VTCS is involved, any RTDs moving out of the CDS will require a VTCS CONFIG RESET to be run to add the specific RTDs from the VTCS configuration.

3. Eject cartridges out of the affected cells to be removed (optional). If cells to be removed contain cartridges, optionally eject all cartridges from the affected ACS.
4. Vary drives OFFLINE to MVS if removing drives (optional). From the MVS console, issue the following command to see if the new drives are offline to MVS.

```
D U,,,xxx-yyy
```

where *xxx-yyy* is the drive ID range (both real and virtual). If they are not, vary the real drives offline with the MVS command:

```
VARY xxx-yyy OFFLINE
```

5. Vary drives OFFLINE to VTCS if removing RTDs (optional). From the MVS console, issue:

```
V RTD(rtd_id) OFF
```

for the RTDs that reside in the LSMs being unallocated.

6. Modify affected CAPs OFFLINE (optional). Issue the following HSC command to no longer use CAPs that are removed from the partition.

```
MODIFY CAP cap-id OFFLINE
```

7. Vary all ACSs connected to the library OFFLINE to all hosts. All ACSs must be offline to all hosts. Issue the HSC command:

```
VARY ACS acs-id OFFLINE
```

8. Run the SET SLIDRIVS utility if deleting drives (optional).

Delete drives to the changed partition (ACS) by running the SLUADMIN utility or command for SET SLIDRIVS.

```
SET SLIDRIVS(,,,,,,,,,,,,,) +
FORLSMID(lsm-id),FORPANEL(12)
```

Note: If VTCS is involved, any RTDs moving out of the CDS will require a VTCS CONFIG RESET to be run to remove the specific RTDs from the VTCS configuration.

9. Change the Library configuration.

The internal configuration file in the SL3000 must be changed to reconfigure the ACS. This change is made through the SLConsole.

- CAPs:

- If removing a CAP, use the SLConsole to unassign the CAP from all partitions.
- If unassigning a CAP, verify the CAP is in manual mode. Use the SLConsole to unassign the CAP from one or more partitions.

- Drives – If removing or unassigning a drive, use the SLConsole to unassign the drive from the partition.

- Drive Bays – If removing a drive bay, use the SLConsole to unassign the drive slots from their partition.

- Cells – If unassigning cells, use the SLConsole to unassign cells from the partition.

- Cartridges left in unassigned cells will become orphaned.

- SL3000:

- Dual Tallbots:

If removing a second robot, no additional setup is needed.

Outermost PEMs must be transformed into CEMs when removing the second robot.

Outermost PEMs must be transformed into CEMs when removing the second robot.

- Expansion Modules:

ATTENTION: Prior to removing a DEM, cartridges must be either moved to available space or ejected if space is not available. After the volumes are removed from the DEM, the partitions need to unassign the cell locations.

ATTENTION: Prior to removing a CEM from the library, cartridges must be either moved to available space or ejected if space is not available. After the volumes are removed from the CEM, the partitions need to unassign the cell locations.

ATTENTION: Prior to removing a CEM from a location other than from either end of the library (i.e., between a PEM and either a CEM, DEM or Base Drive Module) the volumes in the PEM must be either moved to available space within the partition or ejected if no space is available. After

the volumes are removed from the PEM, the partitions need to unassign the cell locations.

10. Vary all ACSs connected to the library ONLINE to all hosts. Issue the HSC command:

```
Vary ACS acs-id Online
```

The HSC retains the online/offline status of each pre-existing LSM. If the LSMs do not modify online when the ACS is varied online, issue the HSC command:

```
MODIFY LSM lsm-id ONLINE
```

No additional mounts will occur to the ACS until the ACS is varied online. All other hosts will complete their configuration process when the ACS is varied online.

Note:

- Virtual tape environments will be required to vary any RTDs online to VTCS. Additionally, any virtual tape drives varied offline will need to be varied online to MVS.
 - Refer to "[Varying the ACS Online After a Configuration Change](#)" for further details of this process.
-
-

11. Re-synchronize the SMC. Issue the SMC RESYNCHRONIZE command.

If you are running multiple HSC hosts, the new configuration will be propagated automatically to all hosts.

Changing the CAP State for Bulk CAP - HSC Down

This procedure details how you can add a Bulk CAP to a partition or remove a Bulk CAP from a partition while the HSC is down.

1. Using the SLC, add the CAP(s) to the partition(s) or remove the CAP(s) from the partition(s) and then apply the changes. See the *SL8500 User's Guide* for information.
2. Start HSC and wait for initialization to complete. See "[Starting and Stopping HSC.](#)" During startup, reconfiguration starts. Wait for reconfiguration to complete.
3. If the LSMs affected are offline:
 - a. Enter `MODify LSM (lsm-range)ONline`
where *lsm-range* is a range of LSMs where the CAP changes are being made.
 - b. Enter `Display Cap` to verify the CAP changes to the library.
 - c. Start HSC activity.
4. If the LSMs affected are online:
 - a. Enter `Display Cap` to verify the CAP changes to the library.
 - b. Start HSC activity.

Changing the CAP State for Bulk CAP - HSC Up

This procedure details how you can add a Bulk CAP to a partition or remove a Bulk CAP from a partition while the HSC is up.

1. Enter `MODify LSM (lsm-range) OFFline`
where *lsm-range* is the range of LSMs where the CAP changes are being made.
2. Stop all HSC activity.
3. Using the SLC, add the CAP(s) to the partition(s) or remove the CAP(s) from the partition(s) and then apply the changes. See the *SL8500 User's Guide* for information.

Reconfiguration starts. Wait for reconfiguration to complete.
4. Enter `MODify LSM (lsm-range) ONline`
where *lsm-range* is the range of LSMs where the CAP changes are being made.
5. Enter `Display Cap` to verify the CAP changes to the library.
6. Start HSC activity.

Varying the ACS Online After a Configuration Change

In several partitioning procedures, there is a “vary ACS online” step that needs to be performed when a hardware configuration change is completed for an SL3000 library. When this step is executed, the HSC will run a reconfiguration for the ACS that is affected. During Reconfiguration of an ACS, the following occurs:

- The ACS configuration containing details of the LSM is read from the library.
- CAP and drive configuration information are read from the library.
- Library configuration information (LSM, CAP, and drive) is then synchronized with the equivalent information in the HSC database.
- LSM cell geography is discovered and updated in the HSC database.

Below is an HSC log example of an SL3000 ACS that change from a library with a Basic Expansion Module (BEM) to a library with a BEM, a DEM, and 2 CEMs:

```
/SLS0000I V ACS 00 ONLINE
/SLS4459I Library configuration changes started
/SLS4400I HSC dynamic reconfiguration initiated
/SLS0668I LSM 00:00 READY; TEMP OUTAGE QUEUE WILL BE RE-DRIVEN.
/SLS0054I ACS 00 now ONLINE
/SLS4415I Configuration change for LSM 00:00, total cells increased by
      1,148, free cells increased by 748
/SLS4419I CAP 00:00:04 status changed from UNALLOCATED to OFFLINE
/SLS4419I CAP 00:00:05 status changed from UNALLOCATED to OFFLINE
/SLS4419I CAP 00:00:07 status changed from UNALLOCATED to OFFLINE
/SLS1969I User exit 06 module SLSUX06 INACTIVE
/SLS4407I HSC dynamic reconfiguration processing has ended.
/SLS4460I Library configuration changes complete on ECCL
```


Configuring VTCS

If you license VTCS, you must configure VTCS with the VTCS CONFIG utility. Refer to the *ELS Command, Control Statement, and Utility Reference* for information about this utility.

Much of the CONFIG process involves defining to VTCS the essential virtual/hardware parts of the VSM solution: VTSSs, VTVs, VTDs, and RTDs. What you also have the option of doing is setting non-default operating values, which is discussed in [Chapter 2, "Planning VTCS Operating Policies."](#)

Tip: ELS no longer requires you to define VTVs and MVCs using CONFIG. Now, it is done through POOLPARAM/VOLPARAM statements, which you created in "[Defining MVCs](#)" and "[Defining VTVs](#)."

This chapter includes two major CONFIG sections:

- "[Building a Simple CONFIG Deck](#)" is basically a procedure where you build a sample basic CONFIG deck step-by-step in an ordered fashion, as described in [Table 6-1](#).
- "[Special Cases: Advanced Uses of CONFIG](#)" shows variations of the CONFIG deck beyond the basic model, such as the deck where you disable a host from initiating migrations.

Table 6-1 VTCS Configuration Tasks

Step	Notes	Planning Information
1	<p>Are you doing an upgrade install? Then you want to run DECOMP to get a true picture of your current configuration. You can then update the DECOMP listing and resubmit it to CONFIG to update your configuration.</p> <p>Typically, you run DECOMP with the highest level of the VTCS (6.0 or 6.1) that supports the CDS "from" level.</p> <p>Or is this a first time install? If so, skip to Step 2.</p>	NA
2	<p>The actual CDS level specification is pretty easy but deciding which level takes some thought, so please study the planning section carefully before you begin coding.</p>	See " CDS VTCS Level Requirements ."

Table 6–1 (Cont.) VTCS Configuration Tasks

Step	Notes	Planning Information
3	Specify global operating policies, such as the number of free MVCs. Remember, you can always take the defaults, and change things later.	"Planning VTCS Operating Policies."
4	Specify the MVC reclamation policies.	"MVC Policies."
5	Specify CONFIG statements for the system's VTSSs, the RTDs attached to the VTSSs, and the VTDs in each VTSS. These statements complete the hardware definitions.	<ul style="list-style-type: none"> ▪ "Planning for VTSSs" ▪ "Planning for Library-Attached Transports" ▪ "Planning for VTDs"

Building a Simple CONFIG Deck

The following procedure shows how to build a simple CONFIG deck. For more complex examples, see "[Special Cases: Advanced Uses of CONFIG.](#)"

Caution: Oracle strongly recommends that you do not "pre-define" hardware through the CONFIG deck (for example, by defining VTSSs that are not physically installed but will be in the future). Only use CONFIG to define hardware that is actually installed.

To Build a Simple CONFIG Deck:

1. Run the VTCS DECOMP Utility if you are performing an upgrade install. Otherwise, skip to Step 2.

The following example shows JCL to run the DECOMP utility with output to flat file CFG22202.

```
//DECOM EXEC PGM=SLUADMIN, PARM='MIXED'
//STEPLIB DD DSN=hlq.SEALINK, DISP=SHR
//SLSCNTL DD DSN=hlq.DBASEPRM, DISP=SHR
//SLSCNTL2 DD DSN=hlq.DBASESEC, DISP=SHR
//SLSSTBY DD DSN=hlq.DBASETBY, DISP=SHR
//CFG22202 DD DSN=hlq.CFG22202, DISP=SHR
//SLSPRINT DD SYSOUT=*
//SLSIN DD *
DECOM FLATDD(CFG22202)
```

2. The following example shows CONFIG JCL to specify CDS level G.

```
//CREATECF EXEC PGM=SLUADMIN, PARM='MIXED'
//STEPLIB DD DSN=hlq.SEALINK, DISP=SHR
//SLSCNTL DD DSN=hlq.DBASEPRM, DISP=SHR
//SLSCNTL2 DD DSN=hlq.DBASESEC, DISP=SHR
//SLSSTBY DD DSN=hlq.DBASETBY, DISP=SHR
//SLSPRINT //SLSIN DD *DD SYSOUT=*
CONFIG CDSLEVEL(V62ABOVE)
```

Note: In this example, for VTCS NCO support, you do not have to specify RESET. RESET is not required when going from F to G level but all hosts accessing the CDS must be at VTCS/NCS 7.0 or above during the CDS level change. Refer to the *ELS Command, Control Statement, and Utility Reference* for details about when RESET is required.

3. Specify global values.

The following example shows a CONFIG deck containing global values. Note that the example specified a Coupling Facility Structure to hold VTCS Lock Data:

- The Structure must be predefined as shown in "[Defining the Coupling Facility Structure to MVS.](#)"
- You now have to specify RESET and all hosts must be down to implement or remove a VTCS Lock Structure within a Coupling Facility.

```
//CREATECF EXEC PGM=SLUADMIN, PARM='MIXED'
//STEPLIB DD DSN=hlq.SEALINK, DISP=SHR
//SLSCNTL DD DSN=hlq.DBASEPRM, DISP=SHR
//SLSCNTL2 DD DSN=hlq.DBASESEC, DISP=SHR
//SLSSTBY DD DSN=hlq.DBASETBY, DISP=SHR
//SLSPRINT DD SYSOUT=*
//SLSIN DD *
CONFIG RESET CDSLEVEL(V62ABOVE)
GLOBAL MAXVTV=65000 MCVFREE=60 VTVATTR=SCRATCH RECALWER=YES
LOCKSTR=STK_VTCS_LOCKS VTVPAGE=LARGE MAXRTDS=32 FASTMIGR=YES
```

Tip: If a Lock Structure is already defined to VTCS, you can use DECOMP, Display CONFIG, and Display LOCKS to display information about the Lock Structure.

4. Specify reclamation policy values. The following example shows a sample CONFIG deck, now displaying reclamation policy values.

```
//CREATECF EXEC PGM=SLUADMIN, PARM='MIXED'
//STEPLIB DD DSN=hlq.SEALINK, DISP=SHR
//SLSCNTL DD DSN=hlq.DBASEPRM, DISP=SHR
//SLSCNTL2 DD DSN=hlq.DBASESEC, DISP=SHR
//SLSSTBY DD DSN=hlq.DBASETBY, DISP=SHR
//SLSPRINT DD SYSOUT=*
//SLSIN DD *
CONFIG RESET CDSLEVEL(V62ABOVE)
GLOBAL MAXVTV=65000 MCVFREE=60 VTVATTR=SCRATCH RECALWER=YES
LOCKSTR=STK_VTCS_LOCKS VTVPAGE=LARGE MAXRTDS=32 FASTMIGR=YES
RECLAIM THRESHLD=70 MAXMVC=30 START=40 CONMVC=5
```

5. Define VTSSs, RTDs, and VTDs.

The following example shows a sample CONFIG deck that defines VTSSs, the RTDs attached to each VTSS (both VSM4s), and the VTDs in each VTSS. Note that the RTD and VTD definitions immediately follow the VTSS statement. Also note that the CMDCONN and IPCONN parameters define the VTSS to host connection as TCP/IP.

```
//CREATECF EXEC PGM=SLUADMIN, PARM='MIXED'
//STEPLIB DD DSN=hlq.SEALINK, DISP=SHR
//SLSCNTL DD DSN=hlq.DBASEPRM, DISP=SHR
//SLSCNTL2 DD DSN=hlq.DBASESEC, DISP=SHR
```

```
//SLSSTBY DD DSN=h1q.DBASETBY,DISP=SHR
//SLSPRINT DD   SYSOUT=*
//SLSIN DD   *
CONFIG RESET CDSLEVEL(V62ABOVE)
GLOBAL MAXVTV=65000 MVCFREE=60 VTVATTR=SCRATCH RECALWER=YES
LOCKSTR=STK_VTCS_LOCKS VTVPAGE=LARGE MAXRTDS=32 FASTMIGR=YES
RECLAIM THRESHLD=70 MAXMVC=30 START=40 CONMVC=5
VTSS NAME=VSM41 LOW=70 HIGH=80 MAXMIG=8 MINMIG=4 RETAIN=5 +
      CMDCONN=TCPIP IPCONN=(192.80.12.1,192.80.12.2)
RTD  NAME=PR11A00 DEVNO=1A00 CHANIF=0C
RTD  NAME=PR11A01 DEVNO=1A01 CHANIF=0D
RTD  NAME=PR11A02 DEVNO=1A02 CHANIF=0K
RTD  NAME=PR11A03 DEVNO=1A03 CHANIF=0L
RTD  NAME=PR12A08 DEVNO=2A08 CHANIF=1C
RTD  NAME=PR12A09 DEVNO=2A09 CHANIF=1D
RTD  NAME=PR12A0A DEVNO=2A0A CHANIF=1K
RTD  NAME=PR12A0B DEVNO=2A0B CHANIF=1L
VTD  LOW=9900 HIGH=99FF
VTSS NAME=VSM42 LOW=70 HIGH=80 MAXMIG=8 MINMIG=4 RETAIN=5 +
      CMDCONN=TCPIP IPCONN=(192.80.12.3,192.80.12.4)
RTD  NAME=PR23A00 DEVNO=3A00 CHANIF=0C
RTD  NAME=PR23A01 DEVNO=3A01 CHANIF=0D
RTD  NAME=PR23A02 DEVNO=3A02 CHANIF=0K
RTD  NAME=PR23A03 DEVNO=3A03 CHANIF=0L
RTD  NAME=PR24A08 DEVNO=4A08 CHANIF=1C
RTD  NAME=PR24A09 DEVNO=4A09 CHANIF=1D
RTD  NAME=PR24A0A DEVNO=4A0A CHANIF=1K
RTD  NAME=PR24A0B DEVNO=4A0B CHANIF=1L
VTD  LOW=9900 HIGH=99FF
```

Updating the Tape Management System for VSM

To update your tape management system (such as CA-1, CA-Dynam/TLMS, and DFSMSrmm), do the following:

- Add volser ranges for VTVs to your tape management system. Ensure that you do not assign vault codes to VTVs.
- Access to the MVCs through an RTD bypasses the MVS intercepts put in place by the tape management system so that it does not record within its database any access to the MVCs by VSM and does not automatically provide protection against inadvertent overwrites of non-expired data on MVCs. Therefore, if you choose to define MVCs to the tape management system, Oracle **strongly recommends** that you define them as non-scratch, non-expiring volumes.
- The tape management system requires an entry in the MVS Subsystem Name Table; this entry must precede the entry for SMC.

If you are using AutoMedia for MVS, ensure that VTVs are defined as virtual volumes to direct AutoMedia to bypass DSN checking, which allows AutoMedia to recall, mount, and reuse non-resident scratch VTVs.

Caution: Note the following:

- VTCS has an automatic interface to notify RMM when a VTV becomes scratch but RMM does not notify VTCS when you unscratch a VTV by changing the CV status to non-scratch. If this is done, you must also run the HSC SLUADMIN utility to unscratch the VTV for VTCS. Otherwise, you may encounter a mount failure when VTCS attempts to select the VTV to service a scratch mount request.
- RMM (DFSMS/RMM) has additional integrity checks at mount time to ensure that the correct volume has been mounted. Because VTCS has features and optimizations that sometimes present a new initialized version of a VTV rather than the current copy of a VTV, it is necessary for VTCS to override these RMM integrity checks. VTCS does these overrides using the LISTVOLUME and CHANGEVOLUME API calls to update the RMM database. You must therefore ensure that HSC has been given the appropriate security access to the RMM API. For more information, refer to your RMM documentation.

Defining MVC Pool Volser Authority

When VSM needs to mount an MVC and to write to an MVC, a SAF query is issued to verify that the HSC user (see ["Defining A Security System User ID for HSC and VTCS"](#)) has UPDATE authority for the MVC. The SAF query is issued on behalf of HSC and passed to the system security product (such as RACF, CA-ACF2, or CA-Top Secret).

VSM requires UPDATE authority for the volsers in the MVC pool. All other users should have an access of NONE for these volsers. Similarly, VSM should not have UPDATE authority for any volsers that are not in the MVC pool. Refer to the documentation for your security product for procedures to add the appropriate tape volume security for VSM. [Table 6-2](#) summarizes these definitions.

Table 6-2 Security Class, Resource Class, and Access Values for MVC Pool Volser Authority

Class	Resource Name	Recommended User Access Levels
Tape volume security	MVC Pool Volume Serials	UPDATE - allows VSM to write on MVC

The following example shows a sample RACF profile and permissions commands to give the user ID VSM8HSC update access to MVC volser CVC024.

```
*****
* Define a profile in the TAPEVOL security class for MVC CVC024 *
*****
RDEFINE TAPEVOL security CVC024 UACC(NONE)
*****
* Allow user ID VSM8HSC update access to MVC CVC024 *
*****
PERMIT CVC024 CLASS(TAPEVOL) ACCESS(UPDATE) ID(VSM8HSC)
*****
```

Caution: Note the following:

- To ensure that MVCs are not accidentally overwritten, for each MVC volser, you must update your TAPEVOL security as described above and your tape management system. For more information, see ["Updating the Tape Management System for VSM."](#)
 - You must also run the HSC UNSCratch Utility to unscratch any current scratch cartridges in the MVC range. For more information, refer to *Managing HSC and VTCS*.
 - Depending on the default settings of your security system, VSM may not be able to mount and to write to MVCs until you have defined a security system user ID for HSC and TAPEVOL profiles for the MVCs.
 - If you add new ranges of MVCs to your VSM system, remember to update the TAPEVOL profiles to include the new ranges.
-
-

Storing VTCS Locks in a Coupling Facility (Optional)

Before VTCS updates a CDS record (for example, a VTV record) it locks the record to avoid contention from concurrent updates from multiple hosts. VTCS releases the CDS lock record once the CDS record has been updated.

Customers with large VTCS configurations experience high CDS I/O rates, part of which is due to the need to access CDS lock records. As configurations grow in size, (for example, by adding more hosts), CDS performance becomes a bottleneck.

An MVS Coupling Facility is a suitable alternative medium for VTCS Lock data because:

- A Coupling Facility provides very fast data transfer speeds, so the new I/O to the Coupling Facility with Lock records is less than the corresponding I/O to a CDS with Lock records.
- MVS provides a technique that allows data stored in a failing Coupling Facility to be re-built in another Coupling Facility (if one exists) without terminating the application.

Note: If VTCS locks are held in a coupling facility structure, VTCS uses the Structure (rather than the HSC mechanism) for sending/receiving Host-to-Host messages.

When to Implement VTCS Locks

Storing VTCS lock data in a Coupling Facility Structure is a solution to the specific problem of VTCS causing high I/O demand to the CDS in some configurations.

Note:

- VTCS locks in a Coupling Facility Structure **only** works if all your hosts are within the same sysplex. Using SMC client/server feature is another solution to the problem that does **not** require all hosts to be within the same sysplex. For more information, refer to *Configuring and Managing SMC*.
- Using a Coupling Facility Structure is **not** a solution to all CDS performance issues. For this reason, Oracle recommends that if you believe you have a CDS performance problem, contact Oracle Software support to have the problem analyzed **before** considering implementing VTCS lock data in a Coupling Facility.

Requirements

To store VTCS Locks in a Coupling Facility:

- All hosts must have access to the same Coupling Facility. Similarly, if you have an alternate Coupling Facility to rebuild the VTCS Lock Structure, all hosts must have access to that alternate Coupling Facility. All hosts must also be in a Sysplex.

The Coupling Facility implementation currently assumes all hosts run HSC/VTCS. VTCS will not start successfully on a host running HSC/VTCS if a host running HSC only is already active.

- The Coupling Facility Structure must be predefined to MVS before VTCS can use it to store CDS Lock Records. VTCS uses the list form of a Coupling Facility Structure. Display LOCKS shows one of the following VTCS Coupling Facility lock types:

Host Footprint

used to serialize access to the host footprint list.

Host to Host

used to serialize access to a given host to host list.

Lock data

used to serialize access to the VTCS lock data.

Formatting

used to serialize the initial formatting of the structure; also used when rebuilding data.

System

lock is held but is not a lock used by VTCS; assume it is used by MVS.

Sizing the Coupling Facility Structure

A Structure size of 768K should be sufficient for configurations up to 100 VTSSs. If the Structure is sized too small, HSC/VTCS will be unable to connect to the Structure or will be able to connect but will be unable to format all of its data. In both cases, VTCS will terminate.

Defining the Coupling Facility Structure to MVS

The following example shows a sample of an IXCMIAPU job to define a VTCS Lock Structure within a Coupling Facility Resource Manager (CRFM). In this example, note that:

- There are two Coupling Facilities, FACIL01 and FACIL02.
- There is a 768K Structure called STK_VTCS_LOCKS to store VTCS Lock Records.
- Structure STK_VTCS_LOCKS can exist in either Coupling Facility but FACIL01 is preferred over FACIL02. If VTCS starts to store lock data in FACIL01 and FACIL01 then becomes unavailable, VTCS attempts to build the STK_VTCS_LOCKS in FACIL02 to ensure continuous operations.

```
//SYSPRINT DD SYSOUT=*
//SYSIN      DD *
DATA TYPE(CFRM) REPORT(YES)
DEFINE POLICY NAME(POLICY1) REPLACE(YES)
  CF  NAME(FACIL01)
  TYPE(123456)
  MFG(IBM)
  PLANT(02)
  SEQUENCE(123456789012)
  PARTITION(1)
  CPCID(00)
  SIZE(0)
  DUMPSPACE(2000)
  CF  NAME(FACIL02)
  TYPE(123456)
  MFG(IBM)
  PLANT(02)
  SEQUENCE(123456789012)
  PARTITION(2)
  CPCID(00)
  SIZE(1)
  DUMPSPACE(2000)
STRUCTURE  NAME(STK_VTCS_LOCKS)
  SIZE(768)
  PREFLIST(FACIL01, FACIL02)
```

Note: If you define only one Coupling Facility, and it becomes unavailable, VTCS terminates on all hosts but HSC will still be running. If this occurs, do one of the following:

- Fix the Coupling Facility error, then recycle HSC/VTCS on all hosts; you can resume without changing the configuration.
 - If the Coupling Facility cannot be fixed:
 1. Stop HSC on all hosts.
 2. Remove LOCKSTR=*structure_name* from the VTCS configuration, then run the CONFIG utility with the RESET parameter (required to remove LOCKSTR=*structure_name*).
 3. Restart the HSCs. The locks will now be maintained in the CDS.
-

Reverting to Storing VTCS Lock Data in the CDS

To revert to storing VTCS Lock Data in the CDS:

1. Stop HSC on all hosts.
2. Run CONFIG RESET **without** the LOCKSTR parameter so VTCS can store VTCS Locks in the CDS.
3. Restart HSC/VTCS on all hosts.

For more information, refer to the *ELS Command, Control Statement, and Utility Reference*.

Managing Failures/Unavailability of the VTCS lock structure

VTCS supports Structure Rebuild to allow for failures/unavailability of the Structure or the Coupling Facility containing the Structure.

Structure rebuild can be initiated by:

- Operator command (SETXCF START,REBUILD,xxx) for a planned outage of the Structure or the Coupling Facility, and/or
- MVS or VTCS detecting an error in, or failure of, the Structure or the Coupling Facility.

Note that VTCS does **not** support System Managed Duplexing.

If the Structure used by VTCS can only be allocated in one Coupling Facility, VTCS will terminate on all Hosts if the Structure (or the Coupling Facility containing the Structure) fails or becomes unavailable.

If the Structure can be allocated in more than one Coupling Facility, VTCS's Structure Rebuild code will attempt to rebuild the data in an alternate Coupling Facility Structure. VTCS will only terminate if the rebuild fails.

Managing Notification of Resource Contention

If a resource is locked (not available), the request cannot be processed and is retried later. At present, error message SLS6946E is issued to provide information about resources that are not available for a period of ten minutes or more because they are locked by another host or task in the system.

The optional GLOBAL CONFIG LOCKTOUT=nnn parameter specifies a value of 0 or a value between 5 and 240. If a value of 0 is specified, message SLS6946E is not issued when a required resource is locked. A value between 5 and 240 specifies the minimum number of minutes that a resource is locked before message SLS6946E is issued. If the parameter is not specified, the default value is 10 minutes. The LOCKTOUT parameter requires F level CDS (V61ABOVE) and above and ELS 7.0 PTF L1H1665 or ELS 7.1 PTF L1H166E.

Configuring VTCS for 512 VTD Support

VSM7 VTSS supports 512 VTDs with current maintenance.

During VTCS and SMC initialization, any VTDs not explicitly defined in VTCS CONFIG are dynamically discovered. Only 256 VTDs can be defined in the VTCS CONFIG, the remaining VTDs are dynamically discovered.

The CONFIG VTSS CMDCONN parameter controls which path (channel or TCP/IP) VTCS uses to send commands to the VTSS. With CMDCOMM=CHANNEL, the

default, only the VTDs defined in CONFIG VTD statements are used to send commands to the VTSS. Therefore, it is recommended to define as many VTDs (up to 256) as possible to CONFIG. CMDCOMM=TCPIP does not have this requirement.

The z/OS HCD configuration must be used to control which hosts have access to the VTDs that are not explicitly defined in the VTCS CONFIG.

Enabling 512 VTD Support

To configure 512 VTD support:

1. Update the IOCP. See [Appendix K, "VSM 6/VSM 7 Configuration."](#)
2. Update the VSM7 or higher with the microcode that supports 512 VTDs.
3. Update ELS with the current maintenance.
4. Update the CONFIG GLOBAL statement by specifying VTDGT256=YES.
5. If necessary, update the CONFIG VTD control statements.
6. Run the CONFIG utility.
7. Start SMC and HSC.

Note:

- If both SMC and HSC are running with 512 VTD support, then 512 VTD support can be implemented by varying the VTSS offline and then online.
 - The HSC Mount and DISMount commands do not support VTDs that are not defined in the VTCS CONFIG. The SMC Mount and DISMount commands must be used in place of the HSC Mount and DISMount commands for any VTDs that are not specified in the VTCS CONFIG.
-
-

Updating or Disabling 512 VTD Support

If necessary, 512 VTD support can be updated or disabled globally for each VTSS.

To update or disable 512 VTD support:

1. Update the CONFIG GLOBAL statement with VTDGT256=YES or disable support for a specific VTSS update with CONFIG VTSS NOGT256.
2. Vary the MVS addresses for VTDs 256 – 511 offline.
3. Vary the VTSS offline and then online.

Starting and Stopping HSC

If you are a new customer, you have to start HSC. If you are an existing customer doing an upgrade and all of your configuration changes are non-disruptive, you will not have to restart HSC.

Note: SMC now supports the z/OS reusable ASID feature, which requires that you do the following:

- Update DIAGxx member in SYS1.PARMLIB with REUSASID(YES)
 - On the start command add REUSASID=YES: S SMC,REUSASID=YES,
-
-

Starting HSC

To start HSC:

1. Ensure that you have completed ELS installation and configuration as described in *Installing ELS and Configuring HSC and VTCS*.
2. Modify your HSC start-up JCL procedure as required. For more information, see "[Creating and Cataloging the HSC Startup Procedure](#)."
3. Enter the MVS START command to run HSC startup procedure. For example:

```
S SLSO
```

If your CDS configuration contains VTCS configuration data, VTCS starts automatically unless you specify the startup parameter NOVTC (for more information, see "[Creating and Cataloging the HSC Startup Procedure](#)." For more information on startup scenarios, see "[HSC Startup Scenarios](#)."

Note:

- If your SMF options for HSC did not specify the SUBTYPE parameter, HSC issues a message that SMF record subtypes 1 through 6 are being recorded.
- If you did not specify GTF EID and FID parameters on the PARM operand in HSC start procedure, HSC issues a message indicating that default EID and FID values are being used.
- HSC can be started prior to hardware arrival to make sure that it has been installed properly. HSC subsystem comes up with the LMU stations offline. HSC operator commands may be entered but any functions requiring interaction with the ACS hardware result in error messages.
- Alternately, a JCL substitution symbol may be specified in HSC cataloged procedure, and additional parameters passed using the START command using the substitution symbol.
- HSC can be initialized to a full or base service level by issuing the MVS START command. HSC can also be preinitialized or initialized under the MVS Master subsystem (MSTR) through the START command.
- In a multiple-host configuration, start one host at a time. **Do not bring up multiple hosts simultaneously.**

4. Start SMC. For more information, refer to *Configuring and Managing SMC*.
5. If necessary, do post-startup tasks. These are required only when the CDS is first initialized:
 - All LSMs are offline, so vary them online with HSC MODify ONline. For subsequent executions of HSC, the last recorded status of the LSMs is obtained from the CDS.
 - All CAP preferences are zero (never selected), so use the CAPPREF command to set CAP preferences.

HSC Startup Scenarios

Table 7–1 shows examples of the START command options, which are further described in the following sections.

Table 7–1 HSC Initialization

HSC Started Under:	HSC Start Command:	Proc Name Same as IEFSSNxx Entry?	SLSBPRESI in IEFSSNxx or Previous HSC Start With PARM='INIT'?
MSTR and PROC in SYS1.PROCLIB	S SLS0	Yes	No
	S SLS0,SUB=MSTR	Yes	Ignored
	S SLS0,PARM='SSYS(SLS0)', SUB=MSTR	No	Ignored

Table 7-1 (Cont.) HSC Initialization

HSC Started Under:	HSC Start Command:	Proc Name Same as IEFSSNxx Entry?	SLSBPRESI in IEFSSNxx or Previous HSC Start With PARM='INIT'?
JES and PROC in Any PROCLIB	S SLS0	Yes	Yes
	S SLS0,PARM='INIT'	Yes	No
	S SLS0,PARM='SSYS(SLS0)'	No	Yes
	S SLS0,PARM='INIT,SSYS(SLS0)'	No	No
	S SLS0,PARM='SSYS(SLS0)'	No	No

IPLing the System

A system IPL is required after defining HSC as a subsystem using the Subsystem Name Table entry IEFSSNxx for the first time. Once HSC has been defined, you do not need to re-IPL to stop and start HSC. If you are migrating from a previous HSC release and HSC is defined in the IEFSSNxx member, then a system IPL is not required. You must perform a COLD start at the first HSC startup. Subsequent HSC startups should not require a COLD start unless a problem is encountered.

If you are using SLSBPRESI to initialize HSC under JES, see ["Preinitializing HSC as an MVS Subsystem"](#) for additional information about starting HSC.

If you are not using SLSBPRESI, refer to ["Starting HSC with PARM='INIT'"](#) and ["Initializing HSC Under the Master Subsystem"](#) for additional information about starting HSC without SLSBPRESI.

Preinitializing HSC as an MVS Subsystem

HSC must be initialized as an MVS Subsystem. There are two ways to preinitialize HSC:

- specifying SLSBPRESI as HSC subsystem initialization routine in the IEFSSNxx member of the system PARMLIB and IPLing MVS as follows:

```
SUBSYS SUBNAME(SLS0) INITRTN(SLSBPRESI) /* keyword format */
```

Note: SLSBPRESI must reside in SYS1.LINKLIB or its link list. After installation, copy the SLSBPRESI module from HSC SEALINK to a library in the MVS link list.

- ["Starting HSC with PARM='INIT'."](#)

Starting HSC with PARM='INIT'

Instead of using SLSBPRESI for HSC subsystem initialization routine, you can specify PARM='INIT' in a START command for a preinitialization startup followed by a START command for the actual startup of HSC. This eliminates the need for HSC's subsystem Master Scheduler Initialization (MSI) routine. MSI functions are taken over by the main HSC address space initialization routines. The use of any subsystem MSI exit routine imposes a restriction that it must be in a data set that is cataloged in the master catalog, and included in the linklist. Previous HSC versions have distributed this module to install in SYS1.LINKLIB. This preinitialization of HSC is for initializing HSC under the primary Job Entry Subsystem (JES).

For example, to preinitialize HSC:

```
S SLS0, PARM='INIT'
```

To then start HSC:

```
S SLS0
```

Initializing HSC Under the Master Subsystem

If you want to initialize HSC under the Master subsystem (MSTR), consider the following:

- If SLSBPRESI is **not** part of the subsystem definition in IEFSSNxx or a previous initialization of the subsystem occurred using the INIT parameter, then enter the following to start HSC under Master subsystem control:

```
S SLS0, SUB=MSTR
```

- If SLSBPRESI is **not** part of the subsystem definition in IEFSSNxx and the subsystem name is the same as the startup procedure name, then no preinitialization or SUB=MSTR is necessary:

```
S SLS0
```

- If you want to initialize HSC under JES later, you can use the INIT parameter to preinitialize HSC.
- When running HSC under the Master subsystem, JES services are not used. It is also possible to receive duplicate messages in the system log.

Starting HSC As a Subsystem

If the EXEC statement PARM parameter specifies a subsystem, you can start HSC as a subsystem. For more information, see ["EXEC Statement."](#) For example:

```
S SLS0, PARM='SSYS(SLS0)'
```

Specifying the HSC Service Level at Startup

To start HSC at Full Service Level:

```
S SLS0
```

To start HSC at Base Service Level:

```
S SLS0, PARM='BASE'
```

The BASE parameter should be used with other parameters in the START command since these parameters override whatever is specified in HSC START Procedure. A sample START command with parameters is:

```
S SLS0, PARM='BASE E(086) F(23) M(00) SSYS(SLS0)'
```

Starting the HSC Behavior Change

For a partitioned SL8500 and SL3000 libraries, NCO will automatically be invoked at HSC startup. This must be done for HSC to update these libraries' configuration for HSC to operate properly. Below is an example of new messages seen during HSC startup:


```

SLS0054I ACS 00 now ONLINE
SLS0668I LSM 00:00 IS OFFLINE; TEMP OUTAGE QUEUE HAS BEEN PURGED.
SLS0054I ACS 01 now ONLINE
SLS0668I LSM 01:00 IS OFFLINE; TEMP OUTAGE QUEUE HAS BEEN PURGED.
SLS2256I Automatic CAP Service AVAILABLE
SLS1505I HSC service level FULL initialization complete
SLS4459I Library configuration changes started
SLS0668I LSM 01:01 IS OFFLINE; TEMP OUTAGE QUEUE HAS BEEN PURGED.
SLS0668I LSM 01:02 IS OFFLINE; TEMP OUTAGE QUEUE HAS BEEN PURGED.
SLS0668I LSM 01:03 IS OFFLINE; TEMP OUTAGE QUEUE HAS BEEN PURGED.
SLS4400I HSC dynamic reconfiguration initiated
SLS4402I CDS drive record not found for existing transport; id = 00:00:10:00
SLS4402I CDS drive record not found for existing transport; id = 00:00:10:02
SLS4405D CDS / LMU conflicts exist; allow dynamic reconfig to proceed; Reply 'Y'
or 'N'
SLS4419I CAP 00:00:01 status changed from UNALLOCATED to OFFLINE
SLS4419I CAP 00:00:03 status changed from UNALLOCATED to OFFLINE
SLS4419I CAP 00:00:04 status changed from UNALLOCATED to OFFLINE
SLS4414I Initial configuration for LSM 00:00, total cells 2,710, free cells 2,710
SLS4407I HSC dynamic reconfiguration processing has ended.
SLS4460I Library configuration changes complete on ECCL
SLS4459I Library configuration changes started
SLS4400I HSC dynamic reconfiguration initiated
SLS4402I CDS drive record not found for existing transport; id = 01:00:01:00
SLS4402I CDS drive record not found for existing transport; id = 01:00:01:01
SLS4402I CDS drive record not found for existing transport; id = 01:00:01:02
SLS4402I CDS drive record not found for existing transport; id = 01:00:01:03
SLS4405D CDS / LMU conflicts exist; allow dynamic reconfig to proceed; Reply 'Y'
or 'N'
SLS4401I LSM/Rail 01:02 successfully unallocated from current HSC configuration
SLS4401I LSM/Rail 01:03 successfully unallocated from current HSC configuration
SLS4407I HSC dynamic reconfiguration processing has ended.
SLS4460I Library configuration changes complete on ECCL

```

Configuration Mismatches

During HSC initialization, HSC remains active if LSM or panel type configuration mismatches occur between the CDS and LMU. Specifically, these mismatches include different:

- Numbers of LSMs
- Or unknown types of LSMs
- Or unknown panel types in an LSM.

In these cases, the affected ACS(s) are forced offline. HSC continues to support the unaffected ACS(s), and you can correct the configuration problems and bring the affected ACSs back online.

Stopping HSC

You can stop HSC in either of two ways:

- By first quiescing HSC, which you do with the MVS STOP command, where:
 1. HSC waits for pending work to complete. This includes actions against the LMU such as current mounts and dismounts, utilities, active CAPs, operator commands, and station termination(s).
 2. HSC updates the CDS and journals as pending work completes.

3. HSC cleans up resources, terminates, and displays a termination message on the system console.
 - By forcing HSC termination with the MVS Cancel command. If this does not stop HSC, enter another Cancel command. If this does not work, use the MVS FORCE command.

Caution: Oracle does not recommend forced termination of HSC because unpredictable results may occur on startups following forced terminations.

A forced termination of HSC causes all HSC processes to abort immediately and termination occurs. Any of the following can occur:

- Volumes can be left in-transit, however, mounted volumes are not affected by forced HSC termination.
- Outstanding utilities can abend.
- The CDS and journals can lose synchronization.

If you do choose to force HSC termination:

1. Before forcing termination, verify that alternative library paths exist for the active SMC to continue performing allocation influencing and mount automation.
 2. Errant volume records are created for volumes left in the playground, CAPs, or passthru ports. Operator intervention may be required to resolve the status of errant volumes when HSC execution is resumed. Transports may require operator intervention to satisfy mount requests.
-

HSC LIBGEN Macros

This appendix describes the HSC LIBGEN macros.

SLIACS

The SLIACS macro defines the following items:

- the communication paths between a host (or hosts) and an ACS
- the LSMs attached to an ACS
- preallocating space in the CDS for additional SL8500 LSMs.

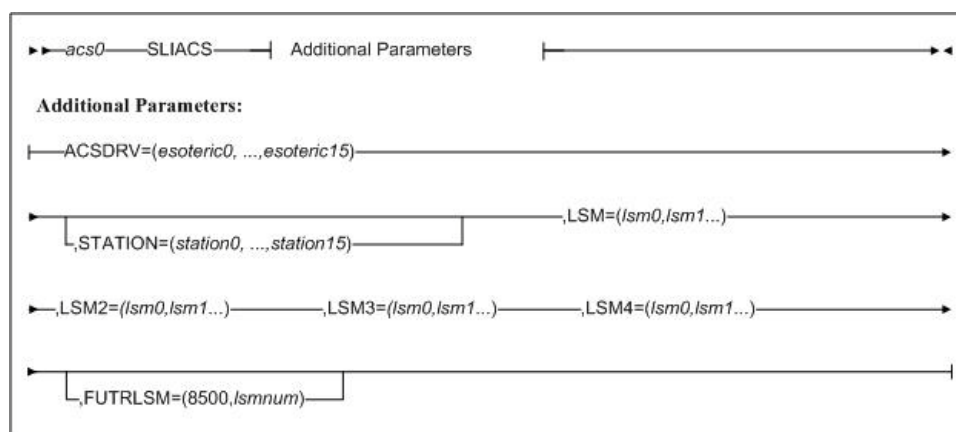
There must be only one SLIACS macro for each ACS.

The SLIACS macro for the first ACS must appear immediately after the SLIALIST macro. The SLIACS macro for each subsequent ACS must appear immediately after the last SLIDRIVS macro for the prior ACS definition.

Syntax

Figure A-1 shows the SLIACS syntax.

Figure A-1 SLIACS Syntax



Parameters

The following information lists the SLIACS parameters.

acs0

assembler label referred to by the SLIALIST macro ACS positional parameter.

SLIACS

name of this macro.

STATION=

station0, ..., station15 are the assembler labels of the SLISTATN macro. They define the station numbers used to communicate between a host and an ACS. Multiple hosts (maximum of 16) may refer to the same SLISTATN macro.

Note: To establish communication with an ACS if the LIBGEN SLISTATN macro has been omitted, use the SET SLISTATN utility to add station addresses for the ACS. Refer to the *ELS Command, Control Statement, and Utility Reference* for a description of this utility.

This parameter is optional.

If STATION is specified, parameters can be omitted for hosts that are not attached to the ACS by using the comma delimiter as follows:

```
STATION=( , STN1)
```

where Host0 does not have a defined connection and STN1 is a defined connection to Host1.

Note: Users running TCP/IP LMU network attachment are not required to supply either SLIACS station number or SLISTATN station address statements. However, if you intend to transition between TCP/IP and 3270 station connections, you should specify both the SLIACS STATION labels and SLISTATN ADDRESS settings.

If STATION is omitted, the ACS has no defined 3270 connections to any host. Leaving out this parameter is desirable if the user wants to define a future ACS and automatically bypass the "ACS acs-id is disconnected" message.

Examples:

If the SLILIBRY macro has been entered as:

```
SLILIBRY HOSTID(MVS0,MVS1
```

and, if the STATION parameter is specified as:

```
STATION=(STN0) ,
```

STN0 corresponds to MVS0.

If STATION is specified as:

```
STATION=(STN0,STN1) ,
```

STN0 corresponds to MVS0, and STN1 corresponds to MVS1.

LSM, LSM2, LSM3, LSM4=

lsm0,lsm1 ... are the assembler labels of the SLILSM macros that define each LSM configured within the ACS.

The number of assembler labels that can be specified on an LSM, LSM2, LSM3, or LSM4 parameter is variable and depends on the length of each label, the number of imbedded commas, and the opening and closing parentheses.

The total number of characters specified cannot exceed the IBM limitation of 255 characters. The LSM2, LSM3, and LSM4 are optional parameters not required if the number of LSMs in the ACS and your naming convention results in a string length less than 255 characters.

For example, if each assembler label is four characters long, a maximum of 50 assembler labels can be specified on the LSM parameter:

200 label characters + 49 imbedded commas + 2 parentheses = 251 characters

Adding another assembler label results in a LIBGEN assembly error. To resolve this problem, add an LSM2 parameter on the SLIACS macro and specify the fifty-first assembler label on that parameter.

In general, if the ACS contains more LSMs (assembler labels) than can be supported by the LSM parameter, you must code LSM2, LSM3, or LSM4 parameters, as necessary, for the other LSMs. It is recommended that you code additional LSMx parameters in ascending order for easier maintainability. Also, note that there is an architectural limit of 256 LSMs for each ACS.

Examples:

```
LSM=(L000)
LSM=(L000,L001,L002)
```

FUTRLSM=

8500 indicates to preallocate space in the CDS for SL8500 LSMs. This space is used when new LSMs are to be added to the ACS through PTP connections.

lsmnum defines the number of future preformatted LSM records written to the CDS. The number specified must be between 4 and 124.

This parameter is optional.

Note:

- An SL8500 library contains four rails (LSMs), thus the *lsmnum* value must be in multiples of four (4, 8, 12, 16, and so on).
- If the number specified in *lsmnum* is too small, another LIBGEN will be required to add additional LSMs. If the number entered for *lsmnum* is too large, the result is unused CDS space reserved for future expansion.


SLIALIST

The SLIALIST macro contains the assembler labels of the SLIACS macro(s). The first ACS listed has an ACSid of 00, the second 01, and so forth.

Syntax

Figure A-2 shows the SLIALIST syntax.

Figure A-2 SLIALIST Syntax



Parameters

acslist

assembler label referred to by the SLILIBRY macro ACSLIST parameter.

SLIALIST

name of this macro.

acs0,acs1,...,acs255

label name used by the SLIACS macro(s). One label name is specified for each ACS, and at least one label name is required.

The order of the SLIACS label names is important. The ACSs are defined in sequence order. For example, label `acs0` assigns that SLIACS macro to ACS 00, label `acs1` assigns that SLIACS macro to ACS 01 and so forth.

If the FUTRACS parameter is specified on the SLILIBRY macro, an unallocated ACS can be created by omitting a label name. For example, to create three ACSs where one ACS is unallocated:

```
ACSLIST SLIALIST ACS0 , , ACS2
```

An unallocated ACS can be added as a functioning ACS to the HSC configuration by using the MODify CONFIG ADD ACSTYPE command as described in "[Adding SL8500 or SL3000 ACSs.](#)"

SLIDLIST

Each SLIDLIST macro specifies an assembler label corresponding to a SLIDRIVS macro which identifies the unique host addresses associated with the transports residing on a cartridge drive panel.

The first SLIDLIST macro for an LSM must appear immediately after the SLILSM macro which refers to it. Subsequent SLIDLIST macros for the LSM must appear immediately after the last SLIDRIVS macro for the preceding SLIDLIST macro.

Syntax

Figure A-3 shows the SLIDLIST syntax.

Figure A-3 SLIDLIST Syntax



Parameters

drvelst0

specifies the assembler label that is referred to by the SLILSM macro DRVELST parameter.

SLIDLIST

name of this macro.

HOSTDRV=

(*drives0,...,drives15*) specifies the assembler label name of each SLIDRIVS macro which defines transport addresses.

Note: The HOSTDRV operands are positional. The SLILIBRY macro HOSTID parameter determines the positional ordering of the operands specified in the HOSTDRV parameter. The first operand specified in the HOSTDRV parameter corresponds to the first host specified in the SLILIBRY macro HOSTID parameter, etc.

Examples:

```
HOSTDRV= (DRV0)
```

where DRV0 is for Host0

```
HOSTDRV= (DRV0 , DRV1)
```

where DRV0 is for Host0 and DRV1 is for Host1

```
HOSTDRV= (DRV0 , DRV0)
```

where DRV0 is for Host0 and DRV0 is for Host1

SLIDRIVS

The SLIDRIVS macro lists the transport device addresses used by a host, which correspond to a specific cartridge drive panel. A SLIDRIVS macro must be specified for each unique operand coded in the SLIDLIST macro HOSTDRV parameter. The SLIDRIVS macro(s) must appear immediately following the SLIDLIST macro which references it.

Syntax

Figure A-4 shows the SLIDRIVS syntax.

Figure A-4 SLIDRIVS Syntax

```
>>-drives0 SLIDRIVS ADDRESS=(addr0,addr1...)-<<
```

Parameters

drives0

specifies the assembler label that is referenced by the SLIDLIST macro HOSTDRV parameter.

SLIDRIVS

name of this macro.

ADDRESS=

addr0,addr1,... specifies the host-unique device addresses associated with transports residing in a cartridge drive panel. Nonexistent transports are indicated by position-holding commas. Duplicate unit addresses are not allowed for a given host.

Note: If addresses exceed more than one line, place a comma after the last address and a nonblank character in column 72 (for example, an X). Continue on the next line starting in column 16. Unlike control statements, no plus (+) or minus (-) continuation characters are required after the last parameter value.

Drives on a panel are defined from top to bottom. When multiple drive columns are present on a 9310 panel, the drive column on the left is defined first, followed by the drive column on the right (as viewed from outside the LSM). The HSC considers these addresses to be attached to an LSM. Verify the addresses you use with your Customer Service Engineer (CSE). [Table A-1](#) shows SLIDRIVS addresses.

Examples:

Table A-1 SLIDRIVS Addresses

ADDRESS=(410,411,412,413)		for a 4480 M24, 4490 M34, 9490 M34, 9490EE M34, or SD-3 H34 device
ADDRESS=(410,,412,)		for a 4480 M22 or 4490 M32 device
ADDRESS=(,B75,,B76)		for a 9490 or 9490EE M32 or SD-3 H32 device
ADDRESS=(,B75,,)		for an SD-3 H31 device
ADDRESS=(,B75,B76,B77)		for an SD-3 H33 device
ADDRESS=(C10,C11,C12,C13,C14, C15,C16,C17,C18,C19)	X	for a 9840, T9840B, T9840C, T9940A, or T9940B device (10-drive panel)
ADDRESS=(C10,C11,C12,C13,C14, C15,C16,C17,C18,C19,C1A,C1B, C1C,C1D,C1E,C1F,C20,C21,C22, C23)	X X X	for a 9840, T9840B, T9840C, T9940A, or T9940B device (20-drive panel)
ADDRESS=(,,,9A01,,,,9A02,,,,9A03,,,9A04)		for an SL8500 device (16-drive panel)
ADDRESS=(,,,00CA,,,,00CB,,,,00CC,,,, 00CD,,,,00CE,,)	X	for an SL3000 device (24-drive panel)

- If you are executing the HSC as a remote server for the SMC, and drive addresses are defined differently on different HSC server hosts, you may need to specify either or both the HSC SET DRVHOST utility or the SMC DRIVEMAP command to define client/server drive mapping. Refer to the SMC DRIVEMAP command in *Configuring and Managing SMC* and the SET DRVHOST utility in the *ELS Command, Control Statement, and Utility Reference*.

- You must specify at least one drive address for each ACS. For example, you cannot specify the following for **all** drive addresses in an ACS.

ADDRESS=(, ,)

Caution: If you fail to enter at least one drive address, the Database Decompile utility will fail against this CDS.

- The total number of drive positions, including position-holding commas, must be 4, 10, 16, or 20 depending on the LSM type:

4410 or 9360 LSM, 4-drive panels are supported.

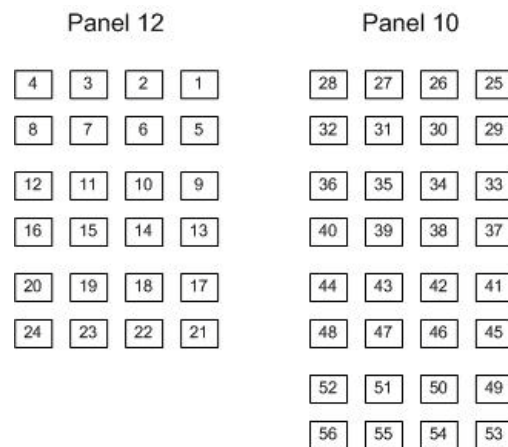
9740 LSM, 4-drive or 10-drive panels are supported.

9310 LSM, 4-drive, 10-drive, and 20-drive panels are supported.

SL3000, 32 drives per panel are supported for panel 10; 24 drives per panel are supported for panel 12. Outside and back of the SL3000 library, drive panel numbering proceeds from top to bottom, right to left. [Figure A-5](#) shows an example of drive numbering for drive addresses on panels 12 and 10.

SL8500, 16-drives per panel are supported.

Figure A-5 SL3000 Drive Panel Layout



- Outside the SL8500 library, drive panel numbering proceeds from top to bottom, right to left, as viewed outside the structure from the drive end. [Table A-2](#) shows an example of drive numbering for drives on one LSM (LSM0) and drive addresses on drives 3 and 15.

Table A-2 SL8500 Drive Panel Layout Example - Outside Library

LSM 0	Drive Panel +2	Drive Panel +1	Drive Panel -1	Drive Panel -2	Row
Rail 1	12	8	4	0	1
Rail 1	13	9	5	1	2
Rail 1	14	10	6	2	3

Table A-2 (Cont.) SL8500 Drive Panel Layout Example - Outside Library

LSM 0	Drive Panel +2	Drive Panel +1	Drive Panel -1	Drive Panel -2	Row
Rail 1	15 (1624)	11	7	3 (1621)	4

- Inside the library, drive panel numbering is top to bottom, left to right, as viewed inside the library from the CAP end. [Table A-3](#) shows the same drive configuration as [Table A-2](#).

Table A-3 SL8500 Drive Panel Layout Example - Outside Library

LSM 0	Drive Panel -2	Drive Panel -1	Drive Panel +1	Drive Panel +2	Row
Rail 1	0	4	8	12	1
Rail 1	1	5	9	13	2
Rail 1	2	6	10	14	3
Rail 1	3 (1621)	7	11	15 (1624)	4

To specify the drives above, the SLIDRIVs statement is:

```
SLIDRIVS ADDRESS=(,,,1621,,,,,,,,,1624)
```

Refer to *Configuring HSC and VTCS* for more information.

- The total number of drive positions specified by all SLIDRIVS statements for a single panel must be the same.
- For an ExtendedStore library, panel 10 must be defined using commas to indicate nonexistent transports. Example:

```
ADDRESS=(,,,) 
```

- On all LSMs except the 9740, drives are defined to the HSC from top to bottom, with addr0 representing the topmost drive and addrn the bottommost drive.

On a 9740 10-drive panel LSM, the drives are populated and configured in the 9740 LSM from bottom to top. (9740 4-drive panels are configured to the 9740 LSM from top to bottom, as are all other LSM drive panels.)

An example showing how to define a 9740 10-drive panel containing five 9840 drives is:

```
ADDRESS=(,,,,,BD4,BD3,BD2,BD1,BD0)
```

SLIENDGN

The SLIENDGN macro specifies the end of the LIBGEN macros. It must appear as the last statement of the LIBGEN. No comments or other statements can follow this macro because they generate assembler warning messages.

Syntax

Figure A-6 shows the SLIENDGN syntax.

Figure A-6 *SLIENDGN Syntax*



Parameters

SLIENDGN

name of this macro. The SLIENDGN macro has no parameters.

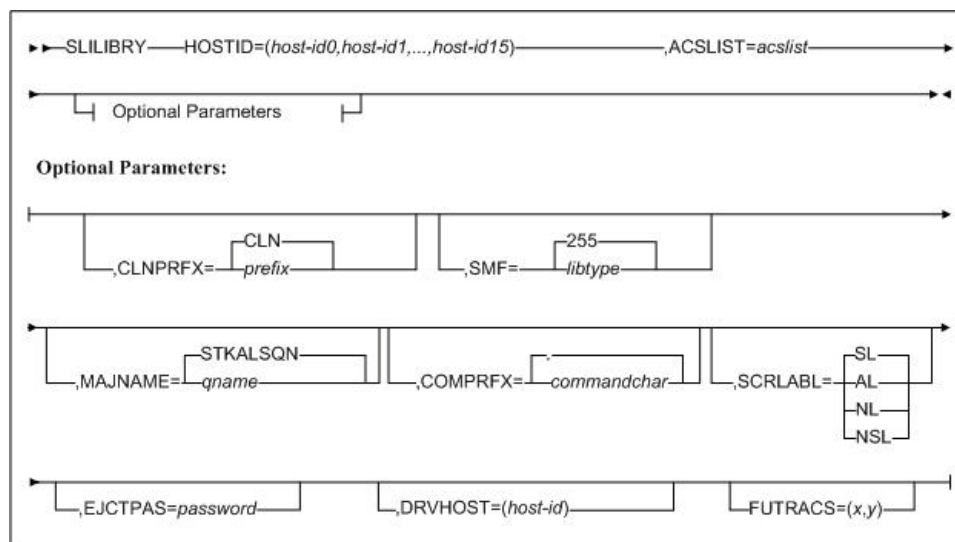
SLIBRARY

The SLIBRARY macro defines the global characteristics of the library. Only one SLIBRARY macro is specified and immediately follows the SLIRCVRY macro.

Syntax

Figure A-7 shows the SLIBRARY syntax.

Figure A-7 *SLIBRARY Syntax*



Parameters

SLIBRARY

name of this macro.

HOSTID=

host-id0, host-id1, ..., host-id15 identifies each host that accesses the library. This represents the SMF system ID for both JES2 and JES3. Valid characters for a HOSTID are A-Z, 0-9, #, \$, and @.

Example:

```
HOSTID= (MVSA, MVSB)
```

where MVSA is Host0 and MVSB is Host1.

Note: The order of the hosts specified is positional. This information is used by other LIBGEN macros. Any changes in host order here may affect other areas of your LIBGEN configuration.

ACSLIST=

acslst is the assembler label of the SLIALIST macro that defines ACSs in the configuration.

Example:

```
ACSLIST=LACSLIST
```

CLNPRFX=

prefix is the three-character prefix that applies to all types of cleaning cartridges. Valid characters are A-Z, 0-9, \$, #, and @. All cartridges having this prefix followed by numerics on their external labels are treated as cleaning cartridges. CLN is the default.

Examples:

```
CLNPRFX=CLN  
CLNPRFX=KCR
```

Note: Cleaning cartridges defined by VOLATTR control statements must still use the cleaning prefix established in the LIBGEN.

SMF=

libtype is the System Management Facility (SMF) record type used by the HSC to record cartridge movement and library performance. The value specified must match the SUBSYS type specified for the HSC in SMFPRMxx. The default is 255.

Example:

```
SMF=255
```

MAJNAME=

qname specifies the ENQ/DEQ/QNAME used by host software for serialization. It must conform to the requirements for an ENQ/RESERVE/DEQ QNAME as defined by the IBM publication *MVS Supervisor Services and Macros*. The default is STKALSQN.

Example:

```
MAJNAME=STKALSQN
```

COMPRFX=

commandchar specifies the command prefix character used to direct operator commands to the HSC. It is a single character. The default command prefix character is a period (.).

Note: Make sure that the prefix character used does not conflict with any of the following:

- another subsystem's command prefix character (such as "\$" for JES2, "*" for JES3, or ";" the delimiter for TSO)
- the command delimiter for MVS — specified as the value for the CMDDELIM keyword in SYS1.PARMLIB (CONSOLxx) and described in the IBM publication *Installation and Tuning Guide for MVS*.
- a JES line editing character as specified in the JES initialization statements or default. For JES2 the initialization statement is CONDEF.

For JES3 the initialization statement is CONSTD. Descriptions of these statements and default values can be found in the *IBM MVS publication Initialization and Tuning Guide*.

- If you are running the HSC on a VM host that shares the CDS with an HSC running on MVS, make sure that your command prefix does not conflict with any of the VM facilities, such as the CP line editing symbol.
-
-

A null character can be specified as the command prefix character. During library operation, to specify an HSC command when the command prefix is a null character, you **must** use the MVS MODIFY command to issue commands to the HSC. [Table A-4](#) shows command prefixes to characters.

Even if you have specified a command prefix, you may use either method for issuing commands to the HSC:

- Specify an HSC command with the command prefix preceding the command.
- Specify an HSC command using the MVS MODIFY command. The following formats for the MVS MODIFY command may be used:

```
MODIFY HSC-subsystem-name, HSC-command
```

or

```
F HSC-subsystem-name, HSC-command
```

where *HSC-subsystem-name* is the HSC subsystem name specified in the HSC startup proc.

Examples of how a specific HSC command can be issued are:

```
MODIFY HSC-subsystem-name, DISPLAY CDS
```

or

```
F HSC-subsystem-name, DISPLAY CDS
```

Table A-4 Mapping Command Prefix Codes to Characters

Hex	Character	Description
40	null	blank
4A	¢	cent
4B	.	period
4C	<	less than
4D	(left parenthesis
4E	+	plus
4F		vertical bar
50	&	ampersand
5A	!	exclamation point
5B	\$	dollar sign
5C	*	asterisk
5D)	right parenthesis
5E	;	semicolon
5F	¬	not symbol
60	-	minus
61	/	slash
6B	,	comma
6C	%	percent
6D	_	underscore
6E	>	greater than
6F	?	question mark
7A	:	colon
7B	#	crosshatch
7C	@	at sign
7E	=	equals
7F	"	double quote

Note: If you specify a null command prefix (hex 40), you must use the MVS MODIFY command to perform any HSC operator command. For example:

```
F SLS0, MOUNT EDU050, B30
```

where:

F is an abbreviation for the MVS MODIFY command

SLS0 is an HSC subsystem

MOUNT is an HSC operator command

EDU050 is a VOLSER

B30 is a designated tape drive

Because of IBM assembler restrictions, a single ampersand (&) cannot be specified as a command prefix. However, if you want to use an ampersand (&) as the command prefix, specify two ampersands (&&). When the LIBGEN file is assembled, the assembler strips off the first ampersand and leaves the second one. The result is that the valid command prefix is a single ampersand.

Example:

```
COMPRFX=&&
```

If you are specifying a left parenthesis (4D) or right parenthesis (5D), the prefix character must be enclosed in single quotes. For example:

```
COMPRFX=' ('
COMPRFX=') '
```

Another example of specifying a valid command prefix is:

```
COMPRFX=@
```

SCRLABL=

specifies the magnetic label type of a library controlled scratch volume. The HSC assumes nonspecific requests with other than the SCRLABL label type are outside the library. If a nonspecific volume is requested with the label type specified, it is considered a scratch volume. Parameter options are:

SL

standard label. The default is SL.

AL

ANSI label

NL

nonlabeled

NSL

nonstandard label.

Automated mounts of scratch volumes other than the SCRLABL label type can be performed by using user exits.

Example:

```
SCRLABL=NSL
```

EJCTPAS=

password specifies that a password is required for the Eject operator command. The password is one to eight alphanumeric characters. Acceptable characters include A-Z (must be capitalized) and numbers 0-9.

Note: Existing passwords do not need to be changed to conform to these restrictions but any new passwords must follow the guidelines described above. If EJCTPAS is not specified, no password is used. An encrypted form of the password is maintained in the control data set.

Example:

```
EJCTPAS=GOODDAY
```

DRVHOST=

hostid specifies a host ID to be used as the "master" host for drive address definitions in a client/server environment. This parameter provides a consistent server environment when a single SMC can communicate with multiple servers that have different defined drive addresses. DRVHOST is used in conjunction with the SMC DRIVEMAP command (refer to Configuring and Managing SMC).

Example:

```
DRVHOST=MVSA
```

FUTRACS=(x,y)

This parameter specifies SL300 or SL8500 ACSs that will be physically installed at a later date. *x* specifies the number of ACSs, and *y* optionally specifies the number of LSMs to add for each added ACS.

Caution: The additional CDS records you create can cause your system to generate an 878-10 ABEND error if there is not enough space to hold the internal records for some operations. If you run out of space during the HSC INIT, you need to re-create your LIBGEN and lower these numbers.

Note:

- The number of ACSs specified can be between 1 and 255 but the total number of real ACSs and FUTRACS must not exceed 255.
 - The number of LSMs for each FUTRACS must be a multiple of 4 and range between 4 and 40.
 - If you do not specify LSMs (*y*), it defaults to 12.
 - The FUTRACSs are generic, meaning you can use them to dynamically add an SL3000 or a string of SL8500s, not to exceed the number of LSMs generated.
 - You **must** specify REGION=6M on your HSC start JCL.
-
-

SLILSM

This macro defines the panel numbers and relative position of cartridge drive panels around an LSM, the assembler labels for the SLIDLIST macro, the panel number for each pass-thru port in an LSM, and the assembler labels of SLILSM macros.

The first SLILSM macro for an ACS must appear immediately after the last SLISTATN macro for an ACS. Subsequent SLILSM macros for the ACS must appear immediately after the last SLIDRIVS macro referred to by the preceding SLILSM macro. The SLILSM macros are coded in the same order as specified in the SLIACS macro LSM, LSM2, LSM3, and LSM4 parameters.

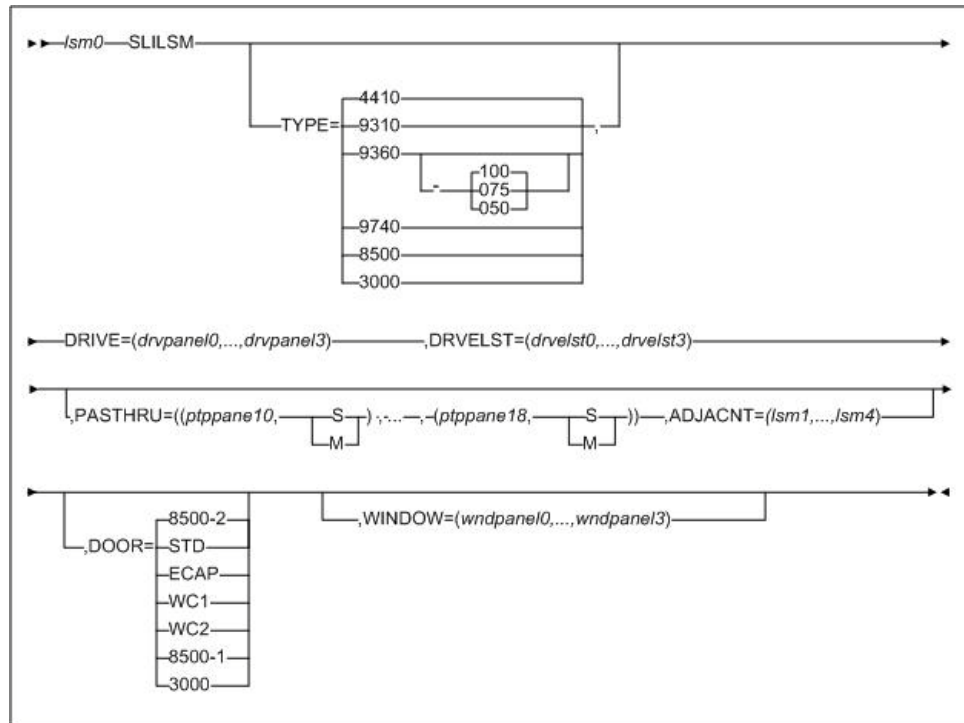
Note:

- For the SL3000 library, the only required parameter is TYPE=3000, although the other parameters are still valid. If you specify only TYPE=3000 and decompile the CDS, the other relevant parameters are shown in the decompile output.
- For the SL8500 library with the enhanced partitioning feature, the only required parameter is TYPE=8500, although the other parameters are still valid. If you specify only TYPE=8500 and decompile the CDS, the other relevant parameters are shown in the decompile output.
- It important to follow the Start from Scratch procedure even if you do not partition the library after completing the LIBGEN. For more information, see "[Starting from Scratch – LIBGEN, SLICREAT.](#)"

Syntax

fix x shows the SLILSM syntax.

Figure A-8 SLILSM Syntax



Parameters

lsm0

assembler label referred to by the SLIACS macro LSM, LSM2, LSM3, and LSM4 parameters.

SLILSM

specifies the name of this macro.

TYPE=

specifies the LSM type.

Caution: The value specified must match both the physical hardware configuration and the value coded by the Oracle CSE or panel mismatch errors appear at startup. Confirm the value you are specifying on the TYPE parameter with your CSE.

Options include:

4410

standard LSM. 4410 is the default.

9310

specifies a PowderHorn LSM.

9360-xxx

specifies a WolfCreek™ LSM with a distinct cartridge capacity.

Allowable values for xxx are:

050

500 cartridge capacity WolfCreek. This is a WolfCreek LSM without cell arrays on panels 3 and 4.

075

750 cartridge capacity WolfCreek. This is a WolfCreek LSM without cell arrays on panel 4.

100

1000 cartridge capacity WolfCreek. This is the default value.

9740

specifies a TimberWolf LSM.

8500

specifies an SL8500 library.

3000

specifies an SL3000 library.

Note: For the SL3000 library, the DRIVE and DRVELST parameters are optional. If you do not specify them, you must run the SET SLIDRIVS utility to add the MVS drive unit addresses.

DRIVE=

dropanel0,...,dropanel3 specifies the range of panel numbers. *dropanel0* specifies the panel number of the first cartridge drive panel moving clockwise from the door; *dropanel1* specifies the second cartridge drive panel, etc. The following are LSM-specific panel requirements:

- For the 4410 standard LSM and the 9310 PowderHorn LSM, *dropanel* must be in the range from 1 through 10, with a maximum of four panels specified. For the ExtendedStore™ LSM, panel 10 must be defined as a drive panel because the SLIDRIVS macro requires a definition for panel 10 (see "SLIDRIVS" for additional information).
- 9360 (WolfCreek) LSMs must have one drive defined in panel 1, with the optional drive/viewing window available in panel 3.
- In 9740 (TimberWolf) LSMs, panel 1 is the drive panel (up to 10 drives can be specified), and panel 3 includes a viewing window if optional cells are not requested.
- For SL8500 libraries, panel 1 is the only drive panel.
- For SL3000 libraries, panels 10 and 12 are the only drive panels.

Examples:

```
DRIVE=(9,10)
DRIVE=(1,2,9,10)
DRIVE=(7,8,9,10)
DRIVE=(1) (WolfCreek or TimberWolf DRIVE example)
```

DRVELST=

drvelst0,...,drvelst3 defines the assembler labels for the SLIDLIST macros. The positional ordering of the DRVELST parameter is determined by the DRIVE parameter.

Example:

```
DRVELST=(PANEL1, PANEL2, PANEL9, PANEL10)
```

PASTHRU=

ptppanel0 specifies the panel number for each pass-thru port in the LSM.

- For 4410 Standard, and 9310 PowderHorn LSMs, *ptppanel0* is a decimal number in the range from 1 through 9 and identifies the panel number of a pass-thru port (PTP).
- For the 9360 WolfCreek and 9740 TimberWolf LSMs, the only valid values for *ptppanel0* are 0 and 2.
- For SL8500, valid values for *ptppanel0* are 0 (internal elevator) and 1 (external PTPs).

Note: The SL8500 contains 3 internal PTPs (elevator) that move cartridges between LSMs (rails).

- For SL3000, this parameter is invalid.

S

indicates this LSM does not control the PTP (slave).

M


indicates this LSM controls the PTP (master).

There can be a total of four pass-thru port panels per 4410 and 9310 LSMs and a total of two pass-thru ports for the 9360, 9740, and SL8500 LSMs.

A maximum of two master (M) pass-thru ports are allowed per LSM as shown in Figure A-9.

For 9360 WolfCreek LSMs, panel zero (0) is always a master PTP. When a 9360 WolfCreek LSM is connected to a 4410 or 9310 LSM, the panel of the 9360 (panel 0 and or 2) connected to the 4410 or 9310 LSM is always a master. When 9360s are connected in a series with other 9360s, panel 0 is always a master and panel 2 is always a slave.

Figure A-9 Example of Pass-thru Port Relationships



G20060

For 9740 TimberWolf LSMs, panel 2 is the master PTP and panel 0 is the slave. 9740 LSMs can attach only to other 9740s.

For SL8500 libraries, the lowest numbered LSM is always the master PTP.

If the PASTHRU parameter is specified, the ADJACNT parameter must also be specified. The order of PTPs listed must correspond to that specified in the ADJACNT parameter.

Examples:

```
PASTHRU=( (5,M) )
PASTHRU=( (5,M) , (7,S) )
PASTHRU=( (1,M) , (3,M) , (5,S) )
PASTHRU=( (0,M) , (2,S) ) (WolfCreek Example)
PASTHRU=( (2,M) , (0,S) ) (9740 Example)
PASTHRU=( (0,M) , (0,M) , (0,M) ) (SL8500 Example)
```

Note: All of parentheses delimiters must be included in the macro statement even if only one PTP is specified. For example, if only one PTP is specified, the double parentheses must be indicated as shown in the following example.

```
PASTHRU=((6,M))
```

ADJACNT=

lsm0,lsm1,lsm2,lsm3 specifies the assembler labels of SLILSM macros, as coded in the SLILSM macro, which are connected through PTPs to this LSM. If the ADJACNT parameter is specified, the PASTHRU parameter must also be specified. The order listed must correspond to that specified in the PASTHRU parameter.

Examples:

ADJACNT= (LSM1)

where LSM1 is *ptppanel0*

ADJACNT= (LSM2 , LSM0)

where LSM2 is *ptppanel0* and LSM0 is *ptppanel2*

Note: For the SL3000 library, this parameter is invalid.

DOOR=

optionally specifies the CAP configuration in the LSM access door.

Note: The 9740 TimberWolf LSM contains either a 10-cell removable magazine or a 14-cell permanent rack. The HSC receives CAP configuration information directly from the LMU, so it is not necessary to specify this parameter for the 9740.

STD

indicates the LSM access door contains a standard CAP (capacity of twenty-one cartridges). Default is STD.

ECAP

indicates the LSM access door contains an enhanced CAP. The enhanced CAP door features two large CAPs (capacity of forty cartridges each), and one small CAP, referred to as the priority CAP or PCAP (capacity of one cartridge).

WC1

indicates the LSM access door contains a WolfCreek CAP having a 20-cell capacity and a PCAP.

WC2

indicates the LSM access door contains an optional WolfCreek CAP having a 30-cell capacity. This is in addition to the capacity provided by the WC1 CAP door. Thus a WolfCreek with an optional WC2 CAP has the following configuration: WC1 with 20-cell capacity and a PCAP plus a WC2 with 30-cell capacity.

8500-1

indicates an SL8500 configured with one CAP (capacity of 39 cartridges).

8500-2

indicates an SL8500 configured with two CAPs (capacity of 39 cartridges each).

3000

indicates an SL3000.

WINDOW=

wndpanel0,...,wndpanel3 specifies one or more panel numbers containing a viewing window. *wndpanel0* specifies the panel number of the first window panel moving clockwise from the access door; *wndpanel1* specifies the second window panel, and so

forth. The panel number(s) designated for the viewing window(s) must be enclosed in parentheses. Viewing windows are an available option on the following LSM types:

- PowderHorn (9310)
- WolfCreek (9360)
- TimberWolf (9740).
- SL3000 - available only on panel 13. If not specified, the HSC self-discovers the window at HSC INIT.

Note: The standard LSM (TYPE=4410) does not allow the replacement of a panel with a viewing window. TYPE=4410 is the default LSM type. You must enter an LSM type of TYPE=9310 (PowderHorn), TYPE=9360-xxx (WolfCreek), or TYPE=9740 (TimberWolf) in order to specify a viewing window.

Only one viewing window can be specified for a WolfCreek or a TimberWolf LSM, and it must be defined in panel number three (3). Selecting any panel number other than three causes an error message to be displayed. If the viewing window is defined for the WolfCreek LSM, panel three may not be designated as a DRIVE panel.

The following example shows the WolfCreek and TimberWolf LSM viewing window selection:

```
WINDOW= (3)
```

The PowderHorn LSM supports from one to four viewing windows. Panel numbers one (1) through nine (9) are eligible for designations as viewing windows. Panel numbers zero (0), ten (10), and eleven (11), are not eligible as viewing windows. If panels zero, ten, or eleven are designated as viewing windows, an error message is issued describing their ineligibility.

The following examples show PowderHorn LSM viewing window selections:

```
WINDOW= (2)
WINDOW= (2, 5, 7, 9)
```

If more than one panel is designated as a viewing window, the panel numbers must be separated by commas and the entire list must be enclosed in parentheses. Panel numbers selected for viewing windows may not be designated as DRIVE panels or as PASTHRU panels.


SLIRCVRY

The SLIRCVRY macro defines criteria to determine when a control data set has become inoperable and requires recovery processing. Only one SLIRCVRY macro is specified, and it is the first macro in the LIBGEN.

Syntax

Figure A-10 shows the SLIRCVRY syntax.

Figure A-10 SLIRCVRY Syntax



Parameters

label

name of the CSECT generated by LIBGEN. If not specified, an unnamed CSECT is generated.

SLIRCVRY

name of this macro.

TCHNIQ

parameter that selects the form of recovery from a CDS failure.

BOTH

specifies that two distinct copies of the control data set (primary and secondary) and journals are specified for recovery purposes. The default is BOTH.

NONE

specifies no form of recovery is used for the control data set. Thus, the primary control data set must be rebuilt, if inaccessible.

SHADOW

specifies that there is to be two distinct copies of the control data set (primary and secondary) for recovery purposes. It is recommended that these data sets reside on separate HDAs and separate strings. A journal is not recorded.

JOURNAL

specifies that there is to be only one primary control data set and that journals are kept. These data sets are to be used for recovery purposes.

The journals contain a record of all transactions that update the control data set. There are two journals per host. It is recommended that they are placed on separate HDAs from the primary control data set DASD volume.

STANDBY

specifies that primary, secondary, and standby control data sets are to be recorded for recovery purposes. No journals are recorded during HSC operation.

ALL

specifies that all control data sets (primary, secondary, and standby) and journals are to be kept and available for recovery purposes.

Note:

- The SLIRCVRY LIBGEN macro TCHNIQE parameter determines how many CDS copies will be initialized by the SLICREAT program and whether or not journals will be initialized by SLICREAT.
- The number of CDS copies used by the HSC is dependent on the number of CDS copies defined in the CDSDEF PARMLIB control statement. It is not determined by the TCHNIQE parameter.
- The HSC uses all of the CDS copies defined in the CDSDEF control statement (whether this includes more or less CDS copies than are specified by the TCHNIQE parameter). However, if journaling is specified by the TCHNIQE parameter, journals must be defined for successful HSC initialization.
- If you are adding a new CDS copy to an existing configuration, you must do the following:
 1. Backup and restore existing CDS copies.
 2. Use SET TCHNIQE to set the recovery technique.

SLISTATN

The SLISTATN macro contains LMU station addresses that connect a host system to an ACS. One SLISTATN macro is used for each station label coded in the SLIACS macro STATION parameter. The SLISTATN macro must appear immediately following the SLIACS macro which references them and in the order specified on the SLIACS macro STATION parameter.

Normally, the HSC waits for offline stations when it attempts to initiate communication. This means message IEF238D is issued when the HSC allocates the station.

However, if the HSC has already established communication with at least one station, the HSC does not wait for other stations to come on during initialization. Therefore, list the stations that are designated for backup usage last. These backup stations are normally offline.

Note: If you are using only TCP/IP LMU connections, you can omit this macro statement but if you are using 3270 or a combination of 3270 with TCP/IP, then you need to include the SLISTATN macro.

Syntax

Figure A–11 shows the SLISTATN syntax.

Figure A–11 *SLISTATN Syntax*

```
▶▶station0——SLISTATN ADDRESS=(addr0,...,addr15)————▶▶
```


Parameters

station0

assembler label referred to by the SLIACS macro STATION parameter.

SLISTATN

name of this macro.

ADDRESS=

addr0,...,addr15 specifies the LMU addresses that connect the ACS to the host system. A minimum of one and a maximum of 16 station addresses may be used by a single host to communicate with an ACS. A maximum of 16 connections per ACS is allowed. In a dual LMU configuration the maximum is 32 connections per ACS; 16 to each LMU in the dual LMU configuration.

Examples:

```
ADDRESS= (0A0)
```

```
ADDRESS= (0A0, 0A1)
```

LIBGEN Macro Samples

This appendix contains the following examples of library configurations and the corresponding LIBGEN files:

- "One Host, Two ACSs, One 9310 LSM, One SL8500 Library Configuration"
- "Four Host, Two ACSs, Two SL8500 Libraries Configuration"
- "Three Host, Three ACSs, Two SL8500 Libraries, One SL3000 Library Configuration"
- "Three Host, Two ACSs, Two SL3000 Libraries Configuration"
- "One Host, One ACS, One 9310 LSM, Dual LMU Configuration"
- "Two Host, One ACS, Two 9310 LSM Configuration"

Note:

- For two host configurations, it is required that host IDs be included in the LIBGEN for each host having access to the library CDS. An example is:

```
SLLIBRY ... HOSTID=(HSC1,HSC2)
```

- HSC **does not distinguish** the type of host; that is, production or test machine.
 - The source code for the sample LIBGENs shown in this appendix is included in the SAMPLIB as members LIBGENnn, where nn is the corresponding LIBGEN example number.
-
-

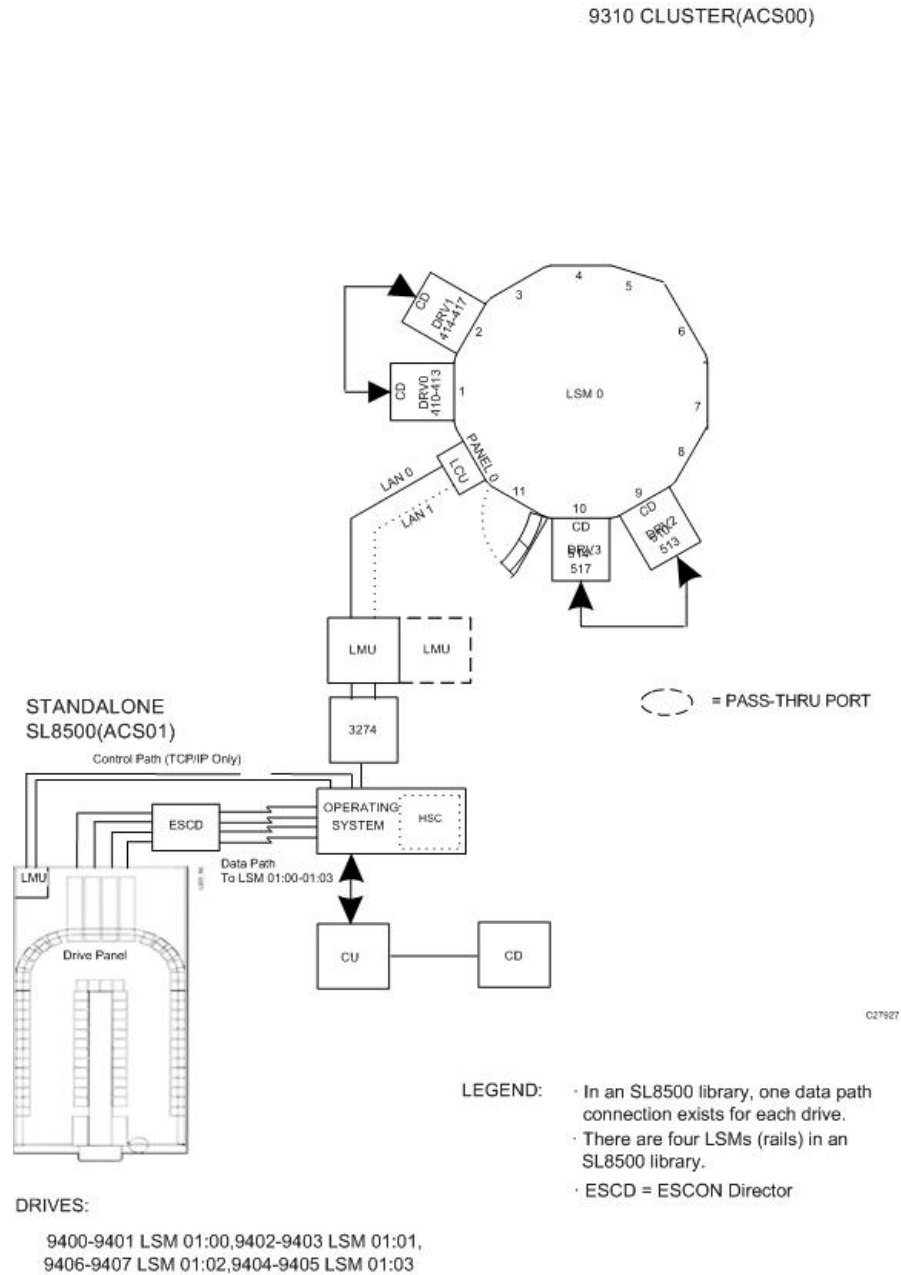
One Host, Two ACSs, One 9310 LSM, One SL8500 Library Configuration

The configuration shown in [Figure B-1](#), which is defined by the LIBGEN01 sample that follows it, consists of the following:

- One host (HSC1)
- Two ACSs (ACS00 and ACS01)
- One 9310 LSM with:
 - 4 transport panels (P000001, P000002, P000009, P000010) and an enhanced CAP
 - Two LMU station addresses (0A0,0A1)
 - 16 transports (410-417 and 510-517)

- One SL8500 standalone library containing 4 LSMs (LSM0100, LSM0101, LSM0102, LSM0103) with:
 - 4 transport panels (P010001, P010101, P010201, P010301)
 - 8 transports (9400-9407).

Figure B-1 One Host, Two ACSs, One 9310 LSM, One SL8500 Library Configuration



* Label definitions:
 *
 * ACS - 'ACSxx' where 'xx' is the ACS number in printable hex.
 * LSM - 'LSMxxyy' where 'xx' is the ACS number in printable hex, and 'yy' is the LSM number within the ACS in printable hex.
 * STATION - 'STxxh' where 'xx' is the ACS number in printable hex, and 'h'


```

D010011 SLIDRIVS ADDRESS=(,,,9400,,,,,9401,,,,,,)
*
LSM0101  SLILSM PASTHRU= ((0,S),(0,S),(0,S)),           X
          ADJACNT=(LSM0100,LSM0102,LSM0103),           X
          DRIVE=(1),                                   X
          DRVELST=(P010101),                           X
          TYPE=8500,                                   X
          DOOR=8500-1
*
P010101  SLIDLIST HOSTDRV=(D0101010,D0101011)
*
D0101010 SLIDRIVS ADDRESS=(,,,,,,,,,9402,,,,,9403)
D0101011 SLIDRIVS ADDRESS=(,,,,,,,,,9402,,,,,9403)
*
LSM0102  SLILSM PASTHRU= ((0,S),(0,S),(0,S)),           X
          ADJACNT=(LSM0100,LSM0101,LSM0103),           X
          DRIVE=(1),                                   X
          DRVELST=(P010201),                           X
          TYPE=8500,                                   X
          DOOR=8500-1
*
P010201  SLIDLIST HOSTDRV=(D0102010,D0102011)
*
D0102010 SLIDRIVS ADDRESS=(,,,,,,,,,9406,,,,,9407)
D0102011 SLIDRIVS ADDRESS=(,,,,,,,,,9406,,,,,9407)
*
LSM0103  SLILSM PASTHRU= ((0,S),(0,S),(0,S)),           X
          ADJACNT=(LSM0100,LSM0101,LSM0102),           X
          DRIVE=(1),                                   X
          DRVELST=(P010301),                           X
          TYPE=8500,                                   X
          DOOR=8500-1
*
P010301  SLIDLIST HOSTDRV=(D0103010,D0103011)
*
D0103010 SLIDRIVS ADDRESS=(,,,9404,,,,,9405,,,,,,)
D0103011 SLIDRIVS ADDRESS=(,,,9404,,,,,9405,,,,,,)
*
SLIENDGN

```

Four Host, Two ACSs, Two SL8500 Libraries Configuration

The configuration, which is defined by the LIBGEN851 sample shown in the following example, consists of the following:

- Four hosts (EC20,EC21,EC22,ECCY)
- Two ACSs (ACS00 and ACS01)
- Two SL8500 libraries connected by PTPs, each containing 4 LSMs with 4 transport panels and transports per panel.

* Label definitions:

*

* ACS- 'ACSxx' where 'xx' is the ACS number in printable hex.

* LSM- 'LSMxxyy' where 'xx' is the ACS number in printable hex, and 'yy' is the LSM number within the ACS in printable hex.

* STATION- 'STxxh' where 'xx' is the ACS number in printable hex, and 'h' is the host index in printable hex.

* PANEL- 'Pxxyypp' where 'xx' is the ACS number in printable hex, 'yy' is the LSM number in printable hex, and 'pp' is the panel

```

* number in decimal.
* DRIVE- 'Dxxyyp'h'where 'xx' is the ACS number in printable hex, 'yy' is
* the LSM number in printable hex, 'pp' is the panel
* number in printable hex, and 'h' is the host index in
* printable hex.
*****
* THIS SAMPLE IS A 2-ACS CONFIGURATION, EACH WITH A SINGLE 8500.
*
LIBGEN851 SLIRCVRY TCHNIQE=SHADOW
*
      SLILIBRY SMF=245,                                X
          ACSLIST=ACSLIST,                            X
          HOSTID=(EC20,EC21,EC22,ECCY) ,             X
          DELDISP=SCRATCH,                            X
          MAJNAME=ARVID,                              X
          CLNPRFX=CLN,                                X
          COMPRFX=-,                                  X
          SCRLABL=SL
*
ACSLIST SLIALIST ACS00,ACS01
*
*   NEXT ACS
*
ACS00   SLIACS LSM=(LSM0000,LSM0001,LSM0002,LSM0003)
*
LSM0000 SLILSM PASTHRU=( (0,M) , (0,M) , (0,M) ) ,    X
          ADJACNT=(LSM0001,LSM0002,LSM0003) ,        X
          DRIVE=(1) ,                                X
          DRVELST=(P00001) ,                          X
          TYPE=8500,                                  X
          DOOR=8500-1
*
P00001 SLIDLIST HOSTDRV=(D000010,D000010,D000010,D000010)
*
D000010 SLIDRIVS ADDRESS=(9000,9001,9002,9003)
*
LSM0001 SLILSM PASTHRU=( (0,S) , (0,M) , (0,M) ) ,    X
          ADJACNT=(LSM0000,LSM0002,LSM0003) ,        X
          DRIVE=(1) ,                                X
          DRVELST=(P00101) ,                          X
          TYPE=8500,                                  X
          DOOR=8500-1
*
P00101 SLIDLIST HOSTDRV=(D001010,D001010,D001010,D001010)
*
D001010 SLIDRIVS ADDRESS=(9010,9011,9012,9013)
*
LSM0002 SLILSM PASTHRU=( (0,S) , (0,S) , (0,M) ) ,    X
          ADJACNT=(LSM0000,LSM0001,LSM0003) ,        X
          DRIVE=(1) ,                                X
          DRVELST=(P00201) ,                          X
          TYPE=8500,                                  X
          DOOR=8500-1
*
P00201 SLIDLIST HOSTDRV=(D002010,D002010,D002010,D002010)
*
D002010 SLIDRIVS ADDRESS=(9020,9021,9022,9023)
*
LSM0003 SLILSM PASTHRU=( (0,S) , (0,S) , (0,S) ) ,    X
          ADJACNT=(LSM0000,LSM0001,LSM0002) ,        X

```

```

                DRIVE= (1) ,                                X
                DRVELST=(P00301) ,                          X
                TYPE=8500,                                    X
                DOOR=8500-1
*
P00301  SLIDLIST HOSTDRV=(D003010,D003010,D003010,D003010)
*
D003010 SLIDRIVS ADDRESS=(9030,9031,9032,9033)
*
* NEXT ACS
*
ACS01   SLIACS LSM=(LSM0100,LSM0101,LSM0102,LSM0103)
*
LSM0100 SLILSM PASTHRU=( (0,M) , (0,M) , (0,M) ) ,          X
                ADJACNT=(LSM0101,LSM0102,LSM0103) ,          X
                DRIVE= (1) ,                                    X
                DRVELST=(P01001) ,                              X
                TYPE=8500,                                       X
                DOOR=8500-1
*
P01001  SLIDLIST HOSTDRV=(D010010,D010010,D010010,D010010)
*
D010010 SLIDRIVS ADDRESS=(9100,9101,9102,9103)
*
LSM0101 SLILSM PASTHRU=( (0,S) , (0,M) , (0,M) ) ,          X
                ADJACNT=(LSM0100,LSM0102,LSM0103) ,          X
                DRIVE= (1) ,                                    X
                DRVELST=(P01101) ,                              X
                TYPE=8500,                                       X
                DOOR=8500-1
*
P00101  SLIDLIST HOSTDRV=(D001010,D001010,D001010,D001010)
*
D001010 SLIDRIVS ADDRESS=(9010,9011,9012,9013)
*
LSM0002 SLILSM PASTHRU=( (0,S) , (0,S) , (0,M) ) ,          X
                ADJACNT=(LSM0000,LSM0001,LSM0003) ,          X
                DRIVE= (1) ,                                    X
                DRVELST=(P00201) ,                              X
                TYPE=8500,                                       X
                DOOR=8500-1
*
P00201  SLIDLIST HOSTDRV=(D002010,D002010,D002010,D002010)
*
D002010 SLIDRIVS ADDRESS=(9020,9021,9022,9023)
*
LSM0003 SLILSM PASTHRU=( (0,S) , (0,S) , (0,S) ) ,          X
                ADJACNT=(LSM0000,LSM0001,LSM0002) ,          X
                DRIVE= (1) ,                                    X
                DRVELST=(P00301) ,                              X
                TYPE=8500,                                       X
                DOOR=8500-1
*
P00301  SLIDLIST HOSTDRV=(D003010,D003010,D003010,D003010)
*
D003010 SLIDRIVS ADDRESS=(9030,9031,9032,9033)
*
* NEXT ACS
*
ACS01   SLIACS LSM=(LSM0100,LSM0101,LSM0102,LSM0103)

```



```

*
LSM0100  SLILSM  PASTHRU=( (0,M) , (0,M) , (0,M) ) ,           X
          ADJACNT=( LSM0101 , LSM0102 , LSM0103 ) ,           X
          DRIVE=( 1 ) ,                                       X
          DRVELST=( P01001 ) ,                                  X
          TYPE=8500 ,                                         X
          DOOR=8500-1

*
P01001   SLIDLIST  HOSTDRV=( D010010 , D010010 , D010010 , D010010 )
*
D010010  SLIDRIVS  ADDRESS=( 9100 , 9101 , 9102 , 9103 )
*
LSM0101  SLILSM  PASTHRU=( (0,S) , (0,M) , (0,M) ) ,           X
          ADJACNT=( LSM0100 , LSM0102 , LSM0103 ) ,           X
          DRIVE=( 1 ) ,                                       X
          DRVELST=( P01101 ) ,                                  X
          TYPE=8500 ,                                         X
          DOOR=8500-1

*
P01101   SLIDLIST  HOSTDRV=( D011010 , D011010 , D011010 , D011010 )
*
D011010  SLIDRIVS  ADDRESS=( 9110 , 9111 , 9112 , 9113 )
*
LSM0102  SLILSM  PASTHRU=( (0,S) , (0,S) , (0,M) ) ,           X
          ADJACNT=( LSM0100 , LSM0101 , LSM0103 ) ,           X
          DRIVE=( 1 ) ,                                       X
          DRVELST=( P01201 ) ,                                  X
          TYPE=8500 ,                                         X
          DOOR=8500-1

*
P01201   SLIDLIST  HOSTDRV=( D012010 , D012010 , D012010 , D012010 )
*
D012010  SLIDRIVS  ADDRESS=( 9120 , 9121 , 9122 , 9123 )
*
LSM0103  SLILSM  PASTHRU=( (0,S) , (0,S) , (0,S) ) ,           X
          ADJACNT=( LSM0100 , LSM0101 , LSM0102 ) ,           X
          DRIVE=( 1 ) ,                                       X
          DRVELST=( P01301 ) ,                                  X
          TYPE=8500 ,                                         X
          DOOR=8500-1

*
P01301   SLIDLIST  HOSTDRV=( D013010 , D013010 , D013010 , D013010 )
*
D013010  SLIDRIVS  ADDRESS=( 9130 , 9131 , 9132 , 9133 )
*
          SLIENDGN

```

Three Host, Three ACSs, Two SL8500 Libraries, One SL3000 Library Configuration

The configuration, which is defined by the LIBGEN853 sample shown in the following example, consists of the following:

- Three hosts (EC20,EC21,EC22)
- Three ACSs (ACS00, ACS01, and ACS02)
- Two SL8500 libraries connected by PTPs, each containing 4 LSMs with 4 transport panels and transports per panel.

■ One SL3000 library whose drives must be defined using SLIDRIVS.

```

* Label definitions:
*
* ACS- 'ACSxx'where 'xx' is the ACS number in printable hex.
* LSM- 'LSMxxyy'where 'xx' is the ACS number in printable hex, and 'yy'
* is the LSM number within the ACS in printable hex.
*STATION- 'STxxh'where 'xx' is the ACS number in printable hex, and 'h'
* is the host index in printable hex.
*PANEL- 'Pxyypp'where 'xx' is the ACS number in printable hex, 'yy' is
* the LSM number in printable hex, and 'pp' is the panel
* number in decimal.
* DRIVE- 'Dxyypph'where 'xx' is the ACS number in printable hex, 'yy' is
* the LSM number in printable hex, 'pp' is the panel
* number in printable hex, and 'h' is the host index in
* printable hex.
*****
* THIS SAMPLE IS A 3-ACS CONFIGURATION WITH TWO SL8500S AND A
* SINGLE SL3000.
*
* NOTE:  FOR THE SL3000, YOU MUST RUN SLIDRIVS TO ADD THE
* MVS UNIT ADDRESSES FOR THE TAPE DRIVES.
*
LIBGEN853 SLIRCVRY TCHNIQE=SHADOW
*
      SLILIBRY SMF=245,                                X
      ACSLIST=ACSLIST,                                X
      HOSTID=(EC20,EC21,EC22),                        X
      DELDISP=SCRTCH,                                X
      MAJNAME=ARVID,                                  X
      CLNPRFX=CLN,                                    X
      COMPRFX=-,                                      X
      SCRLABL=SL
*
ACSLIST  SLIALIST ACS00,ACS01,ACS02
*
*   FIRST ACS
*
ACS00    SLIACS LSM=(LSM0000,LSM0001,LSM0002,LSM0003)
*
LSM0000  SLILSM PASTHRU=((0,M),(0,M),(0,M)),          X
          ADJACNT=(LSM0001,LSM0002,LSM0003),        X
          DRIVE=(1),                                  X
          DRVELST=(P00001),                            X
          TYPE=8500,                                  X
          DOOR=8500-2
*
P00001   SLIDLIST HOSTDRV=(D000010,D000010,D000010)
*
D000010  SLIDRIVS ADDRESS=(9000,9001,9002,9003)
*
LSM0001  SLILSM PASTHRU=((0,S),(0,M),(0,M)),          X
          ADJACNT=(LSM0000,LSM0002,LSM0003),        X
          DRIVE=(1),                                  X
          DRVELST=(P00101),                            X
          TYPE=8500,                                  X
          DOOR=8500-2
*
P00101   SLIDLIST HOSTDRV=(D001010,D001010,D001010)
*
D001010  SLIDRIVS ADDRESS=(9010,9011,9012,9013)

```

```

*
LSM0002  SLILSM PASTHRU= ((0,S),(0,S),(0,M)),           X
          ADJACNT= (LSM0000,LSM0001,LSM0003),           X
          DRIVE= (1),                                    X
          DRVELST= (P00201),                             X
          TYPE=8500,                                     X
          DOOR=8500-2

*
P00201   SLIDLIST HOSTDRV=(D002010,D002010,D002010)
*
D002010  SLIDRIVS ADDRESS=(9020,9021,9022,9023)
*
LSM0003  SLILSM PASTHRU= ((0,S),(0,S),(0,S)),           X
          ADJACNT= (LSM0000,LSM0001,LSM0002),           X
          DRIVE= (1),                                    X
          DRVELST= (P00301),                             X
          TYPE=8500,                                     X
          DOOR=8500-2

*
P00301   SLIDLIST HOSTDRV=(D003010,D003010,D003010)
*
D003010  SLIDRIVS ADDRESS=(9030,9031,9032,9033)
*
*   NEXT ACS
*
ACS01    SLIACS LSM=(LSM010)
*
LSM010   SLILSM TYPE=3000
*
*   NEXT ACS
*
ACS02    SLIACS LSM=(LSM0200,LSM0201,LSM0202,LSM0203)
*
LSM0200  SLILSM PASTHRU= ((0,M),(0,M),(0,M)),           X
          ADJACNT= (LSM0201,LSM0202,LSM0203),           X
          DRIVE= (1),                                    X
          DRVELST= (P02001),                             X
          TYPE=8500,                                     X
          DOOR=8500-2

*
P02001   SLIDLIST HOSTDRV=(D020010,D020010,D020010)
*
D020010  SLIDRIVS ADDRESS=(9200,9201,9202,9203)
*
LSM0201  SLILSM PASTHRU= ((0,S),(0,M),(0,M)),           X
          ADJACNT= (LSM0200,LSM0202,LSM0203),           X
          DRIVE= (1),                                    X
          DRVELST= (P02101),                             X
          TYPE=8500,                                     X
          DOOR=8500-2

*
P02101   SLIDLIST HOSTDRV=(D021010,D021010,D021010)
*
D021010  SLIDRIVS ADDRESS=(9210,9211,9212,9213)
*
LSM0202  SLILSM PASTHRU= ((0,S),(0,S),(0,M)),           X
          ADJACNT= (LSM0200,LSM0201,LSM0203),           X
          DRIVE= (1),                                    X
          DRVELST= (P02201),                             X
          TYPE=8500,                                     X

```

```

                DOOR=8500-2
*
P02201  SLIDLIST HOSTDRV=(D022010,D022010,D022010)
*
D022010  SLIDRIVS ADDRESS=(9220,9221,9222,9223)
*
LSM0203  SLILSM PASTHRU=( (0,S) , (0,S) , (0,S) ) ,                X
          ADJACNT=(LSM0200,LSM0201,LSM0202) ,                    X
          DRIVE=(1) ,                                             X
          DRVELST=(P02301) ,                                       X
          TYPE=8500,                                             X
          DOOR=8500-2
*
P02301  SLIDLIST HOSTDRV=(D023010,D023010,D023010)
*
D023010  SLIDRIVS ADDRESS=(9230,9231,9232,9233)
*
*
                SLIENDGN

```

Three Host, Two ACSs, Two SL3000 Libraries Configuration

The configuration, which is defined by the following example, consists of the following:

- Three hosts (EC20,EC21,9999)
- Two ACSs (ACS00, ACS01)
- Two SL3000 libraries whose drives must be defined using SLIDRIVS.
- The FUTRACS parameter allows you to add two SL8500s or two SL3000s in the future.

```

* Label definitions:
*
* ACS- 'ACSxx'where 'xx' is the ACS number in printable hex.
* LSM- 'LSMxxyy'where 'xx' is the ACS number in printable hex, and 'yy'
* is the LSM number within the ACS in printable hex.
* STATION- 'STxxh'where 'xx' is the ACS number in printable hex, and 'h'
* is the host index in printable hex.
* PANEL- 'Pxxyypp'where 'xx' is the ACS number in printable hex, 'yy' is
* the LSM number in printable hex, and 'pp' is the panel
* number in decimal.
*DRIVE- 'Dxxyypph'where 'xx' is the ACS number in printable hex, 'yy' is
* the LSM number in printable hex, 'pp' is the panel
* number in printable hex, and 'h' is the host index in
* printable hex.
*****
*
* This sample is a 2-ACS configuration, each with a single SL3000.
* It also allows for 2 more SL3000's or 2 SL8500's to be added
* in the future.
*
* NOTE: For the second SL3000, you must run SLIDRIVS to add the
* MVS unit addresses for the tape drives.
*
LIBGEN33 SLIRCVRY TCHNIQE=SHADOW
*
                SLILIBRY SMF=245,                X
                FUTRACS=(2,4) ,                    X

```

```

        ACSLIST=ACSLIST,                                X
        HOSTID=(EC20,EC21,9999),                        X
        MAJNAME=SUNENQ,                                  X
        CLNPRFX=CLN,                                    X
        COMPRFX=-,                                       X
        SCRLABL=SL
*
ACSLIST  SLIALIST ACS00,ACS01
*
* First ACS (SL3000)
*
ACS00    SLIACS LSM=(LSM000)
*
LSM000   SLILSM TYPE=3000,                              X
        DRIVE=(12),                                     X
        DRVELST=(P000012)
*
P000012  SLIDLIST HOSTDRV=(D0000120,D0000120,D0000120)
*
D0000120 SLIDRIVS ADDRESS=(3001,3002,3003,3004,3005,3006,3007,3008,   X
        3009,3010,3011,3012,3013,3014,3015,3016,   X
        3017,3018,3019,3020,3021,3022,3023,3024,   X
        3025,3026,3027,3028,3029,3030,3031,3032)
*
*
* Second ACS (SL3000)
*
ACS01    SLIACS LSM=(LSM001)
*
LSM001   SLILSM TYPE=3000
*
        SLIENDGN

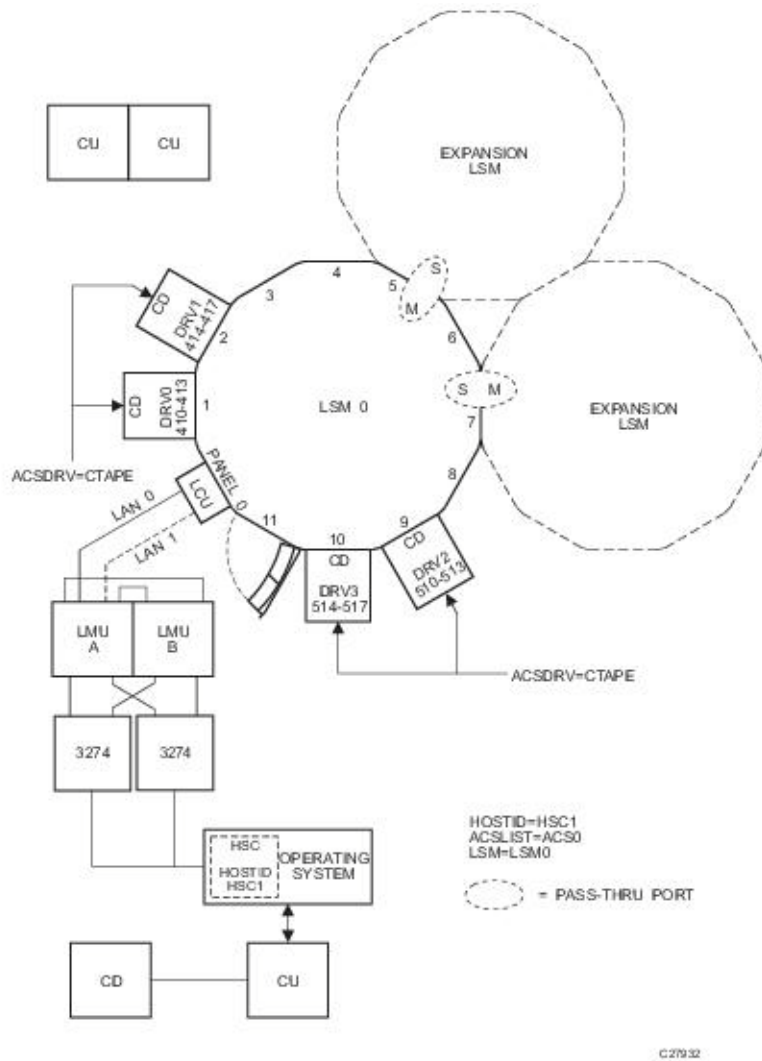
```

One Host, One ACS, One 9310 LSM, Dual LMU Configuration

The configuration shown in [Figure B-2](#), which is defined by the LIBGEN04 sample that follows it, consists of the following:

- One host (HSC1)
- One 9310 LSM (LSM0000)
- 4 LMU station addresses (0A0,0A1,0C0,0C1)
- 4 transport panels (1,2,9,10)
- 16 transports (410-417 and 510-517).

Figure B-2 One Host, One ACS, One 9310 LSM, Dual LMU Configuration



*Label definitions:
 *
 * ACS- 'ACSxx'where 'xx' is the ACS number in printable hex.
 * LSM- 'LSMxxyy'where 'xx' is the ACS number in printable hex, and 'yy'
 * is the LSM number within the ACS in printable hex.
 * STATION- 'STxxh'where 'xx' is the ACS number in printable hex, and 'h'
 * is the host index in printable hex.
 * PANEL- 'Pxxyypp'where 'xx' is the ACS number in printable hex, 'yy' is
 * the LSM number in printable hex, and 'pp' is the panel
 * number in decimal.
 * DRIVE- 'Dxxyypph'where 'xx' is the ACS number in printable hex, 'yy' is
 * the LSM number in printable hex, 'pp' is the panel
 * number in printable hex, and 'h' is the host index in
 * printable hex.

 *
 LIBGEN04 SLIRCVRY TCHNIQE=BOTH
 *
 SLILIBRY HOSTID=HSC1, X
 SMF=235, X

```

                                COMPRFX=#,                                X
                                ACSLIST=NAMEACS,                          X
                                MAJNAME=STKALSQN,                          X
                                SCRLABL=SL,                                X
                                EJCTPAS=GOODDAY,                           X
                                CLNPRFX=CLN,                               X
*
NAMEACS  SLIALIST  ACS00
*
ST000    SLISTATN  ADDRESS=(0A0,0A1,0C0,0C1)
*
LSM0000  SLILSM    DRIVE=(1,2,9,10),                                     X
                                DRVELST=(P000001,P000002,P000009,P000010),  X
                                TYPE=9310,                                   X
                                DOOR=STD
*
P000001  SLIDLIST  HOSTDRV=D0000000
D0000000 SLIDRIVS  ADDRESS=(410,411,412,413)
P000002  SLIDLIST  HOSTDRV=D0000010
D0000010 SLIDRIVS  ADDRESS=(414,415,416,417)
P000009  SLIDLIST  HOSTDRV=D0000020
D0000020 SLIDRIVS  ADDRESS=(510,511,512,513)
P000010  SLIDLIST  HOSTDRV=D0000030
D0000030 SLIDRIVS  ADDRESS=(514,515,516,517)
*
                                SLIENDGN

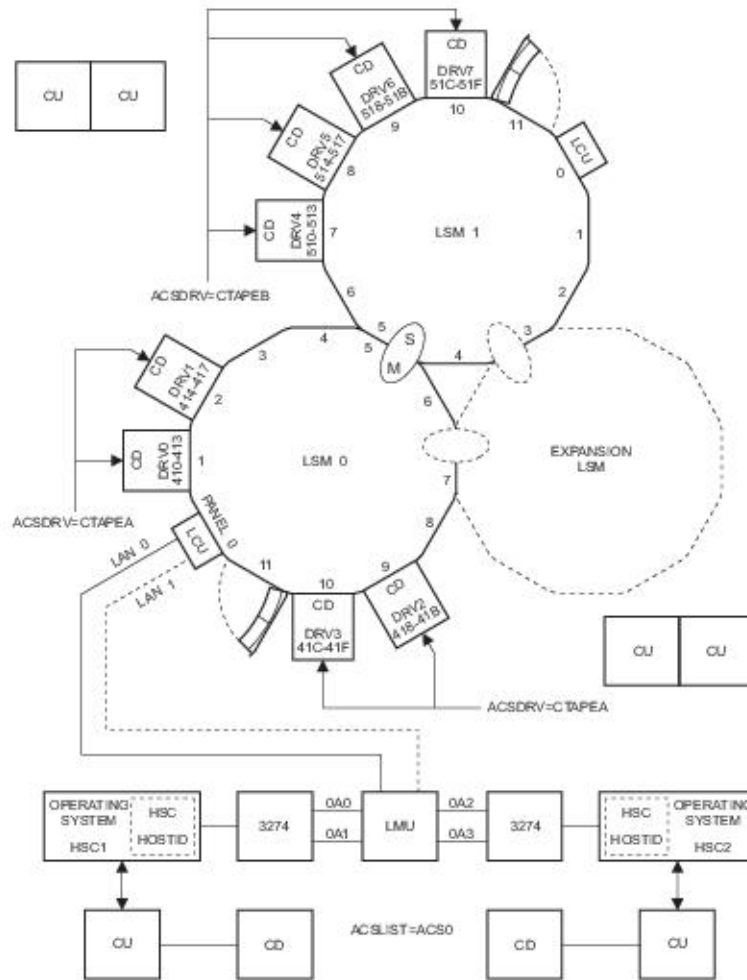
```

Two Host, One ACS, Two 9310 LSM Configuration

The configuration shown in [Figure B-3](#), which is defined by the LIBGEN05 sample that follows it, consists of the following:

- Two hosts (HSC1,HSC2).
- One Automated Cartridge System (ACS00).
- One 9310 LSM designated as LSM0000 with four transport panels (1,2,9,10) and a standard CAP.
- One 9310 LSM designated as LSM0001 with four transport panels (7,8,9,10) and an enhanced CAP.
- 4 LMU station addresses (0A0,0A1,0A2,0A3).
- 32 transports (410-41F and 510-51F).

Figure B-3 Two Host, One ACS, Two 9310 LSM Configuration



C27928

```

* Label definitions:
*
* ACS- 'ACSxx'where 'xx' is the ACS number in printable hex.
* LSM- 'LSMxxyy'where 'xx' is the ACS number in printable hex, and 'yy'
* is the LSM number within the ACS in printable hex.
* STATION- 'STxxh'where 'xx' is the ACS number in printable hex, and 'h'
* is the host index in printable hex.
* PANEL- 'Pxxypp'where 'xx' is the ACS number in printable hex, 'yy' is
* the LSM number in printable hex, and 'pp' is the panel
* number in decimal.
* DRIVE- 'Dxxypph'where 'xx' is the ACS number in printable hex, 'yy' is
* the LSM number in printable hex, 'pp' is the panel
* number in printable hex, and 'h' is the host index in
* printable hex.
*****
LIBGEN05 SLIRCVRY  TCHNIQE=SHADOW
*
      SLILIBRY SMF=235,                               X
              HOSTID=(HSC1,HSC2),                     X
              ACSLIST=NAMEACS,                          X
    
```



```

                                MAJNAME=STKALSQN,                X
                                COMPRFX=#,                       X
                                SCRLABL=SL,                     X
                                EJCTPAS=GOODDAY,                X
                                CLNPRFX=CLN,                   X
*
NAMEACS  SLIALIST  ACS00
*
ST000    SLISTATN  ADDRESS=(0A0,0A1)
ST001    SLISTATN  ADDRESS=(0A2,0A3)
*
LSM0000  SLILSM    DRIVE=(1,2,9,10),                          X
                                DRVELST=(P000001,P000002,P000009,P000010),  X
                                PASTHRU=( (5,M) ),                X
                                ADJACNT=(LSM0001),                X
                                TYPE=9310,                          X
                                DOOR=STD
*
P000001  SLIDLIST  HOSTDRV=(D0000000,D0000000)
D0000000 SLIDRIVS  ADDRESS=(410,411,412,413)
P000002  SLIDLIST  HOSTDRV=(D0000010,D0000010)
D0000010 SLIDRIVS  ADDRESS=(414,415,416,417)
P000009  SLIDLIST  HOSTDRV=(D0000020,D0000020)
D0000020 SLIDRIVS  ADDRESS=(418,419,41A,41B)
P000010  SLIDLIST  HOSTDRV=(D0000030,D0000030)
D0000030 SLIDRIVS  ADDRESS=(41C,41D,41E,41F)
*
LSM0001  SLILSM    DRIVE=(7,8,9,10),                          X
                                DRVELST=(P000107,P000108,P000109,P000110),  X
                                PASTHRU=( (5,S) ),                X
                                ADJACNT=(LSM0000),                X
                                TYPE=9310,                          X
                                DOOR=ECAP
*
P000107  SLIDLIST  HOSTDRV=(D0000040,D0000040)
D0000040 SLIDRIVS  ADDRESS=(510,511,512,513)
P000108  SLIDLIST  HOSTDRV=(D0000050,D0000050)
D0000050 SLIDRIVS  ADDRESS=(514,515,516,517)
P000109  SLIDLIST  HOSTDRV=(D0000060,D0000060)
D0000060 SLIDRIVS  ADDRESS=(518,519,51A,51B)
P000110  SLIDLIST  HOSTDRV=(D0000070,D0000070)
D0000070 SLIDRIVS  ADDRESS=(51C,51D,51E,51F)
*
                                SLIENDGN

```

Special Cases: Advanced Uses of CONFIG

In this section, you are going to learn about the functionality of the CONFIG utility. Note that in many of these examples, you are changing the configuration, which is something you typically do when you are doing an upgrade install: upgrade the hardware configuration, run DECOMP, then update the CONFIG deck to match the hardware changes at a release boundary.

Note that when you are doing something like adding RTDs, there is more to it than just plugging in the hardware and updating the CONFIG deck. You need to do all the tasks discussed in "[Planning for Library-Attached Transports.](#)"

CONFIG Example: All Hosts Access VTDs in One VTSS, Only Selected Hosts Access VTDs in Second VTSS

The following shows sample CONFIG JCL to define a VSM configuration as follows:

- The VTD statement specifies default VTD addresses 8900 - 893F for VTSS1. All hosts have access to these VTDs by their default addresses.
- No default VTD addresses are specified for VTSS2. The VTD statements that immediately follow the HOST statements for MVS1 and MVS2 specify that only these hosts can access the VTDs in VTSS2 by the addresses 9900 - 993F. HOST statement MVS3 is a placeholder; this host cannot access the VTDs in VTSS2, and is disabled from initiating migrates and reclaims.

```
//UPDATECFG EXEC PGM=SLUADMIN, PARM='MIXED'
//STEPLIB DD DSN=hlq.SEALINK, DISP=SHR
//SLSCNTL DD DSN=hlq.DBASEPRM, DISP=SHR
//SLSCNTL2 DD DSN=hlq.DBASESEC, DISP=SHR
//SLSSSTBY DD DSN=hlq.DBASESTBY, DISP=SHR
//CFG22202 DD DSN=hlq.CFG22202, DISP=SHR
//SLSPRINT DD SYSOUT=*
//SLSIN DD *
CONFIG CDSLEVEL(V62ABOVE)
GLOBAL MAXVTV=65000 MVCFREE=60 VTVATTR=SCRATCH RECALWER=YES
LOCKSTR=STK_VTCS_LOCKS VTVPAGE=LARGE MAXRTDS=32 FASTMIGR=YES
RECLAIM THRESHLD=70 MAXMVC=30 START=40 CONMVC=5
VTSS NAME=VTSS1 LOW=70 HIGH=80 MAXMIG=3 RETAIN=5
RTDNAME=VT128800 DEVNO=8800 CHANIF=0A
RTDNAME=VTS18801 DEVNO=8801 CHANIF=0I
RTDNAME=VTS18802 DEVNO=8802 CHANIF=1A
RTDNAME=VTS18803 DEVNO=8803 CHANIF=1I
VTD LOW=8900 HIGH=893F
VTSS NAME=VTSS2 LOW=70 HIGH=80 MAXMIG=3 RETAIN=5
RTDNAME=VTS28804 DEVNO=8804 CHANIF=0A
RTDNAME=VTS28805 DEVNO=8805 CHANIF=0I
```

```

RTDNAME=VTS28806 DEVNO=8806 CHANIF=1A
RTDNAME=VTS28807 DEVNO=8807 CHANIF=1I
HOST    NAME=MVS1
VTDLOW=9900 HIGH=993F
HOST    NAME=MVS2
VTD    LOW=9900 HIGH=993F
HOST    NAME=MVS3 NOMIGRAT NORECLAM

```

CONFIG Example: Update Configuration to Add RTDs

The following example shows JCL to run CONFIG to add RTDs VTS18811 and VTS18813 (connected to VTSS1) to the configuration. Because of NCO, you can add the RTDs dynamically (no RESET required), by simply adding new definitions for RTDs VTS18811 and VTS18813.

```

//UPDATECFG EXEC PGM=SLUADMIN, PARM='MIXED'
//STEPLIB DD DSN=hlq.SEALINK, DISP=SHR
//SLSCNTL DD DSN=hlq.DBASEPRM, DISP=SHR
//SLSCNTL2 DD DSN=hlq.DBASESEC, DISP=SHR
//SLSSTBY DD DSN=hlq.DBASETBY, DISP=SHR
//CFG22202 DD DSN=hlq.CFG22202, DISP=SHR
//SLSPRINT DD  SYSOUT=*
//SLSIN DD  *
CONFIG CDSLEVEL(V62ABOVE)
GLOBAL MAXVTV=65000 MCVFREE=60 VTVATTR=SCRATCH RECALWER=YES
LOCKSTR=STK_VTCS_LOCKS VTVPAGE=LARGE MAXRTDS=32 FASTMIGR=YES
RECLAIM THRESHLD=70 MAXMVC=30 START=40 CONMVC=5
VTSS NAME=VTSS1 LOW=70 HIGH=80 MAXMIG=3 RETAIN=5
RTDNAME=VTS18800 DEVNO=8800 CHANIF=0A
RTDNAME=VTS18801 DEVNO=8801 CHANIF=0I
RTDNAME=VTS18802 DEVNO=8802 CHANIF=1A
RTDNAME=VTS18803 DEVNO=8803 CHANIF=1I
RTDNAME=VTS18811 DEVNO=8811 CHANIF=0E
RTDNAME=VTS18813 DEVNO=8813 CHANIF=1E
VTDLOW=8900 HIGH=893F
VTSS NAME=VTSS2 LOW=70 HIGH=80 MAXMIG=3 RETAIN=5
RTDNAME=VTS28804 DEVNO=8804 CHANIF=0A
RTDNAME=VTS28805 DEVNO=8805 CHANIF=0I
RTDNAME=VTS28806 DEVNO=8806 CHANIF=1A
RTDNAME=VTS28807 DEVNO=8807 CHANIF=1I
HOST    NAME=MVS1
VTD    LOW=9900 HIGH=993F
HOST    NAME=MVS2
VTD    LOW=9900 HIGH=993F
HOST    NAME=MVS3

```

CONFIG Example: Update Configuration to Add MVCs and VTVs and Change AMTs

The following example shows JCL to run CONFIG to modify the configuration shown in the previous JCL example.

- Adding VTVs C25001 to C50000 as scratch.
- Adding MVCs N45000 to N45999.
- Changing the LAMT to 50 and the HAMT to 85 on both VTSS1 and VTSS2.

```

//UPDATECFG EXEC PGM=SLUADMIN, PARM='MIXED'
//STEPLIB DD DSN=hlq.SEALINK, DISP=SHR

```

```

//SLSCNTL DD DSN=hlq.DBASEPRM,DISP=SHR
//SLSCNTL2 DD DSN=hlq.DBASESEC,DISP=SHR
//SLSSTBY DD DSN=hlq.DBASETBY,DISP=SHR
//CFG22202 DD DSN=hlq.CFG22202,DISP=SHR
//SLSPRINT DD  SYSOUT=*
//SLSIN DD  *
CONFIG CDSLEVEL(V62ABOVE)
GLOBAL MAXVTV=65000 MCVFREE=60 VTVATTR=SCRATCH RECALWER=YES
LOCKSTR=STK_VTCS_LOCKS VTVPAGE=LARGE MAXRTDS=32 FASTMIGR=YES
RECLAIM THRESHLD=70 MAXMVC=30 START=40 CONMVC=5
VTVVOL LOW=C25001 HIGH=C50000 SCRATCH
MVCVOL LOW=N45000 HIGH=N45999
VTSS NAME=VTSS1  LOW=50 HIGH=85  MAXMIG=3  RETAIN=5
RTDNAME=VTS18800 DEVNO=8800 CHANIF=0A
RTDNAME=VTS18801 DEVNO=8801 CHANIF=0I
RTDNAME=VTS18802 DEVNO=8802 CHANIF=1A
RTDNAME=VTS18803 DEVNO=8803 CHANIF=1I
RTDNAME=VTS18811 DEVNO=8811 CHANIF=0E
RTDNAME=VTS18813 DEVNO=8813 CHANIF=1E
VTDLOW=8900 HIGH=893F
VTSS NAME=VTSS2  LOW=50 HIGH=85  MAXMIG=3  RETAIN=5
RTDNAME=VTS28804 DEVNO=8804 CHANIF=0A
RTDNAME=VTS28805 DEVNO=8805 CHANIF=0I
RTDNAME=VTS28806 DEVNO=8806 CHANIF=1A
RTDNAME=VTS28807 DEVNO=8807 CHANIF=1I
HOST  NAME=MVS1
VTD  LOW=9900 HIGH=993F
HOST  NAME=MVS2
VTD  LOW=9900 HIGH=993F
HOST  NAME=MVS3

```

CONFIG Example: Denying Host Access to a Physically Removed VTSS

If you remove a VTSS, you need a way to let the hosts know the VTSS is no longer in the configuration. The following example shows JCL to run CONFIG to deny host access to VTSS2 that you physically removed from your configuration. In this example, you simply respecify the VTSS statement for VTSS2 with no parameters to deny host access to this VTSS.

```

//UPDATECFG EXEC PGM=SLUADMIN,PARM='MIXED'
//STEPLIB DD DSN=hlq.SEALINK,DISP=SHR
//SLSCNTL DD DSN=hlq.DBASEPRM,DISP=SHR
//SLSCNTL2 DD DSN=hlq.DBASESEC,DISP=SHR
//SLSSTBY DD DSN=hlq.DBASETBY,DISP=SHR
//CFG22202 DD DSN=hlq.CFG22202,DISP=SHR
//SLSPRINT DD  SYSOUT=*
//SLSIN DD  *
CONFIG CDSLEVEL(V62ABOVE)
GLOBAL MAXVTV=65000 MCVFREE=60 VTVATTR=SCRATCH RECALWER=YES
LOCKSTR=STK_VTCS_LOCKS VTVPAGE=LARGE MAXRTDS=32 FASTMIGR=YES
RECLAIM THRESHLD=70 MAXMVC=30 START=40 CONMVC=5
VTSS NAME=VTSS1  LOW=70 HIGH=80  MAXMIG=3  RETAIN=5
RTDNAME=VTS18800 DEVNO=8800 CHANIF=0A
RTDNAME=VTS18801 DEVNO=8801 CHANIF=0I
RTDNAME=VTS18802 DEVNO=8802 CHANIF=1A
RTDNAME=VTS18803 DEVNO=8803 CHANIF=1I
RTDNAME=VTS18811 DEVNO=8811 CHANIF=0E
RTDNAME=VTS18813 DEVNO=8813 CHANIF=1E
VTDLOW=8900 HIGH=893F
VTSS NAME=VTSS2

```

CONFIG Examples: Using the Maximum 32 RTDs Feature

As described in "[Maximum RTDs per VTSS](#)," you can specify up to 32 RTDs attached to a VTSS. With this feature enabled, it turns out you can also attach greater than 16 RTDs and have some ports left for CLINK originators and terminators. What follows is a general procedure for implementing Maximum 32 RTDs with links to two examples with CONFIG decks.

Implementing Support for Maximum of 32 RTDs

1. Ensure that your system has the Synchronous Replication requirements described in [Table H-3](#).
2. Use CONFIG GLOBAL to enable support for maximum of 32 RTDs.

```
CONFIG GLOBAL MAXRTDS=32
```

Note: Enabling support for a maximum of 32 RTDs does not require CONFIG RESET. However, regressing from 32 RTDs supported to 16 RTDs supported does require CONFIG RESET.

3. Update your CONFIG RTD and CONFIG CLINK statements as required.

Configuring SL8500 Library Communications

The SL8500 is different from other LSMs, and the way it communicates with HSC is also different, as explained in this appendix.

TCP/IP Communications - Important Considerations

The following topics are discussed here:

- ["SL8500 TCP/IP Connections"](#)
- ["Shared Networks"](#)

SL8500 TCP/IP Connections

When you create an SL8500 library complex by connecting libraries together with passthru ports, all hosts must connect to **only one** library in the complex— preferably to the first or rightmost library.

For the complex, there are three types of TCP/IP connections for SL8500 discussed in the following sections:

- ["Dual IP Connections"](#) - A redundancy feature to allow up to two connections to the first library of a complex.
- ["Multiple SL8500 Library Connections"](#) - Starting with the SL8500 3.9x firmware release, an SL8500 complex is capable of network connecting up to 4 libraries to HSC. This enhances connecting redundancy to the complex compared to dual TCP/IP.
- ["Multiple TCP/IP Redundant Electronics"](#) - Starting with the SL8500 6.x firmware release, an SL8500 complex is capable of redundant electronics of the library controller with network connecting up to 16 libraries in pairs (one active and one standby) to HSC. This enhances hardware reliability to the complex compared to the other 2 types.

Note: It is crucial that you **read completely** these sections to understand these procedures.

Shared Networks

The following are some examples of issues that can arise when you connect an SL8500 to a shared network.

- A TCP/IP-connected library can handle standard host traffic but it cannot resolve floods of Address Resolution Protocol (ARP) broadcasts. **For this reason, it is best to attach the library on a controlled network, such as behind a switch or router.**

Later generation networks, such as 1000Base-T and Gig-E, support earlier communication modes; however, devices that are communicating with the library may transmit data at bandwidths that could overwhelm the library.

It is best to attach the library on a controlled network, such as with a switch that can isolate the library from network broadcasts.

- When you connect the library on shared networks and broadcasts are sent to all network nodes, they may also be directed to the library (even though it does not need them).

During the time the library is receiving these irrelevant broadcasts, it cannot receive from or respond to other requests in a timely fashion. This heavy broadcast traffic on the network can saturate the library to the point that, to the host, it may appear that the TCP/IP connection has been lost. It may also result in Overdue Response Handler (ORH) messages on the host system.

- Heavy network traffic can also overwhelm the Ethernet controller causing the processor to continuously reset and re-initialize the controller, then recover the host-to-library communications.

Dual IP Connections

The LMUPATH control statement allows users to define network LMU attachments (refer to the *ELS Command, Control Statement, and Utility Reference*). In a dual IP connection environment for an SL8500, specify a second LMUADDR parameter to define dual IP. HSC automatically determines whether or not the connection is dual IP or dual LMU.

Note: Vary the ACS offline and back online to pick up the revised LMUPATH statement that includes the second connection.

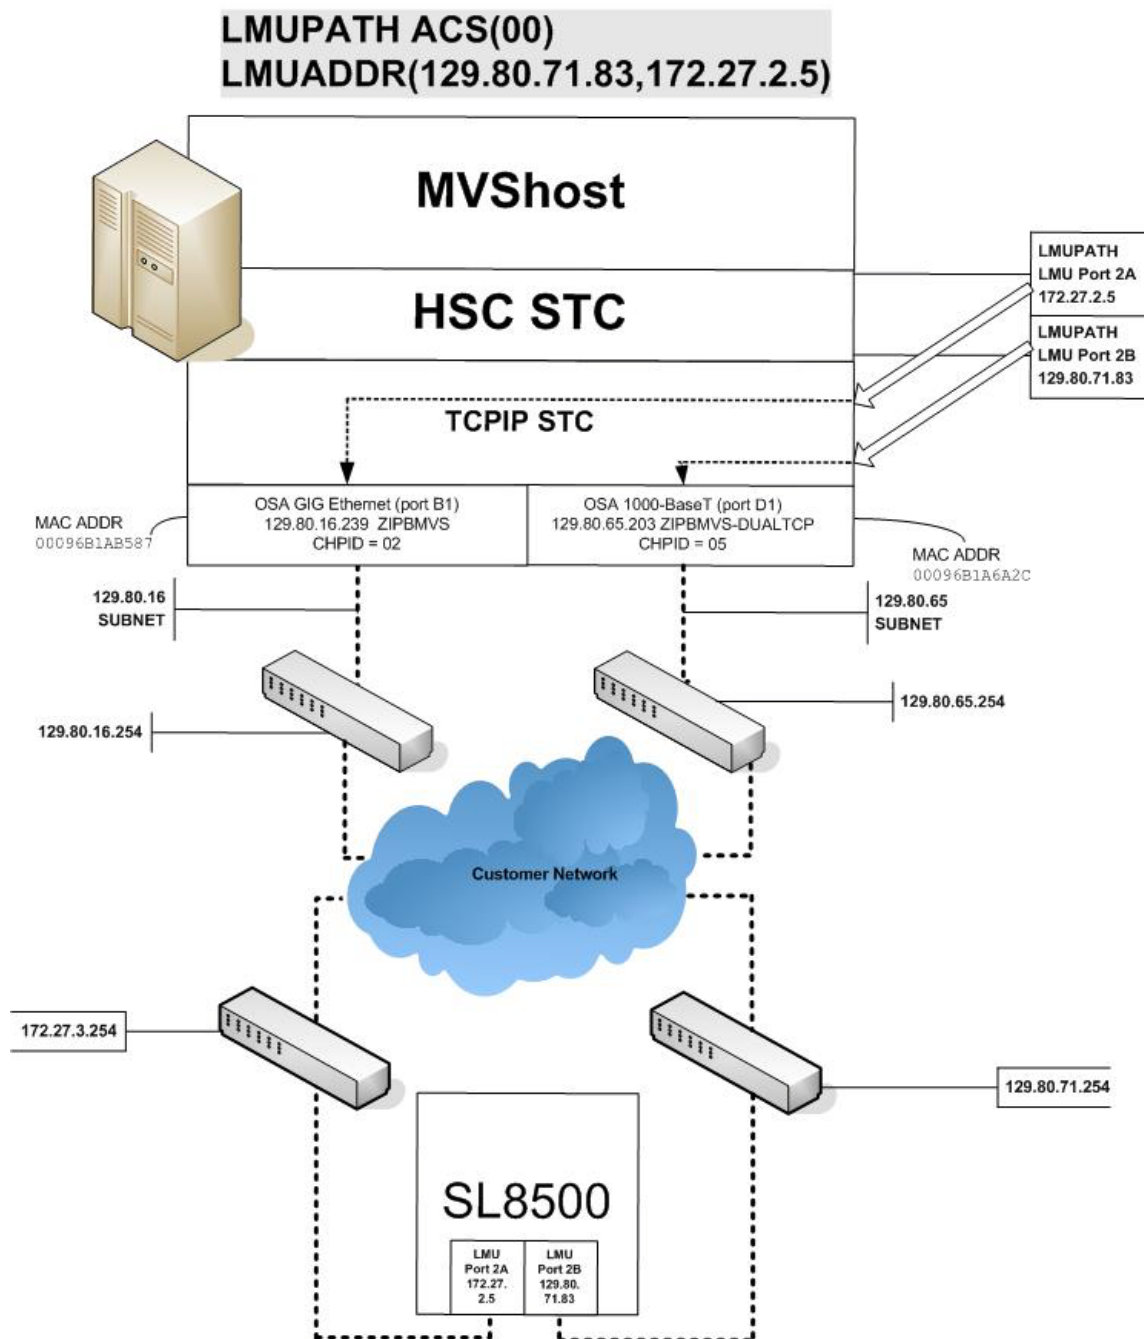
[Figure D-1](#) and [Figure D-2](#) show examples of a dual IP connection for the HSC.

Two SL8500 Network Connections - Two Mainframe IP Addresses

To establish two dedicated routes on different subnetworks from the HSC to two dedicated routes to the SL8500, use the process described in this section. To configure the SL8500 dual TCP/IP feature, refer to *StreamLine SL8500 Modular Library System Dual TCP/IP Feature* or *StreamLine SL3000 Modular Library System Installation Guide*.

[Figure D-1](#) shows dual IP connections with two SL8500 network connections.

Figure D-1 Dual IP Connection with two SL8500 Connections



1. Complete the Network Entries Worksheet (see [Table D-1](#) for a sample) for each set of two dedicated routes to the SL8500. The worksheet can be found in the StreamLine SL8500 Modular Library System Dual TCP/IP Feature document.

Table D-1 Network Entries Worksheet

Port Number	Description	IP Address
2A	HBC Card Port 2A	172.27.2.5
2A	Gateway Port 2A	172.27.3.254

Table D-1 (Cont.) Network Entries Worksheet

Port Number	Description	IP Address
2A	Netmask	/23
2A	Host to Port 2A	129.80.16.239
2B	HBC Card Port 2B	129.80.71.83
2B	Gateway Port 2B	129.80.71.254
2B	Netmask	/23
2B	Host to Port 2B	129.80.65.203

- Define a second DEVICE and LINK statement in your TCP/IP profile data set for a second mainframe network connection.

Example:

```
; OSA CARD #1
DEVICE ECCQD01 MPCIPA NONROUTER AUTORESTART
LINK ZIPBMVS IPAQENET ECCQD01
```

```
; OSA CARD #2
DEVICE ECCQA01 MPCIPA NONROUTER AUTORESTART
LINK ZIPB2MVS IPAQENET ECCQA01
```

- Define a second home address in your TCP/IP profile data set.

Example:

```
HOME
129.80.16.239 ZIPBMVS
129.80.65.203 ZIPB2MVS
```

- Define a second router on the second subnetwork in the routing paragraph of your TCP/IP profile data set. Remember that you must also configure the SL8500 routing tables according to the instructions described in the *StreamLine SL8500 Modular Library System Dual TCP/IP Feature* document.

Example:

```
BEGINROUTES
; NETWORK MASK FIRSTHOP LINKNAME PACKETSIZE
ROUTE 129.80.16.0/24 = ZIPBMVS MTU 1492
ROUTE 129.80.65.0/24 = ZIPB2MVS MTU 1492
ROUTE 172.27.2.5 HOST 129.80.16.254 ZIPBMVS MTU 1492
ROUTE 129.80.71.83 HOST 129.80.65.254 ZIPB2MVS MTU 1492
ROUTE DEFAULT 129.80.16.254 ZIPBMVS MTU 1492
ROUTE DEFAULT 129.80.65.254 ZIPB2MVS MTU 1492
ENDROUTES
```

- Define two dedicated static routes to the SL8500 destination port (2A and 2B) IP addresses over two different routers.

Example:

```
; NETWORK MASK ROUTER LINKNAME PACKETSIZE
ROUTE S18500-port-2A-IP-Address HOST 129.80.16.254 MVSHOST1 MTU 1492
ROUTE S18500-port-2B-IP-Address HOST 129.80.64.254 MVSHOST2 MTU 1492
```

```

BEGINROUTES
;
NETWORK MASK    FIRSTHOP    LINKNAME    PACKETSIZE
ROUTE 129.80.16.0/24    =    ZIPBMVS    MTU 1492
ROUTE 129.80.65.0/24    =    ZIPB2MVS    MTU 1492
ROUTE 172.27.2.5    HOST    129.80.16.254    ZIPBMVS    MTU 1492
ROUTE 129.80.71.83    HOST    129.80.65.254    ZIPB2MVS    MTU 1492
ROUTE DEFAULT 129.80.16.254    ZIPBMVS    MTU 1492
ROUTE DEFAULT 129.80.65.254    ZIPB2MVS    MTU 1492
ENDROUTES

```

6. Start the second mainframe network connection device.

```
V TCPIP, tcp-stc-name, START, device_name
```

7. Define a second LMUADDR parameter for the port 2A IP address on the SL8500.

Example:

```
LMUPATH ACS(00) LMUADDR(129.80.71.83,172.27.2.5)
```

8. Enter the LMUPDEF command containing the LMUPATH statements that define the host name or IP address for each ACS.

```
LMUPDEF DSN('xxx.xxx.xxx(xxx)')
```

In the following example, LMUPDEF loads LMUPATH parameters from YOUR.DSN(MEMBER)

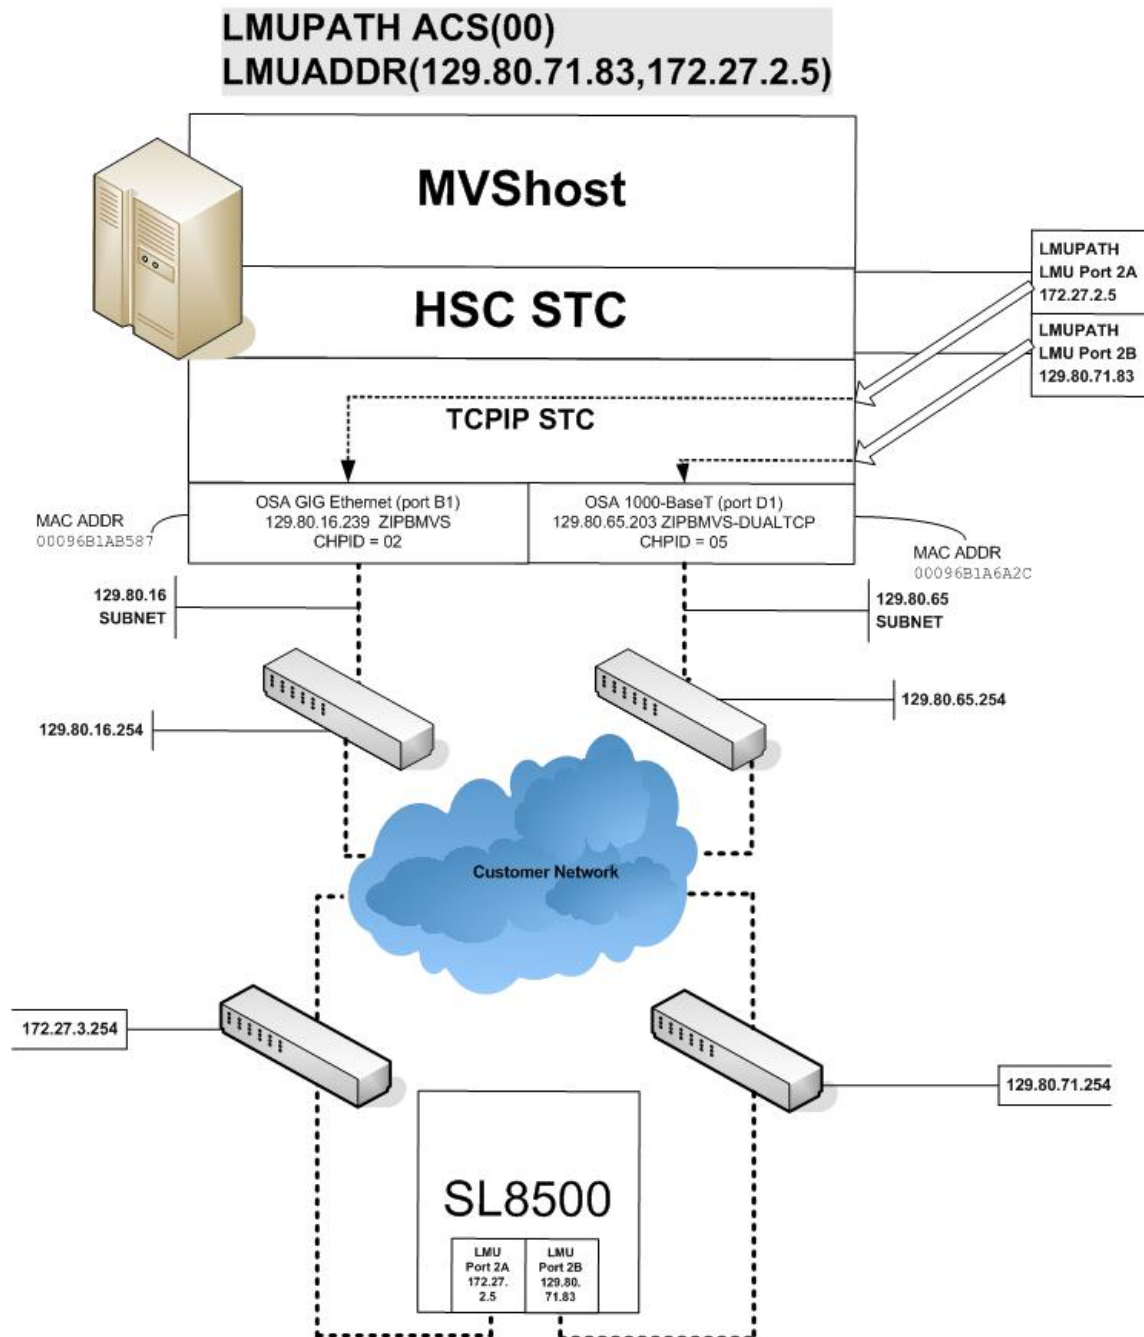
```
LMUPDEF DSN('YOUR.DSN(MEMBER)')
```

9. Allow the trained SL8500 service representative to enter the network connections to the SL8500 library for either port 2A and 2B, whichever is applicable.
10. Vary the ACS offline and back online to pick up the revised LMUPATH statement that includes the second connection. This can be done one host at a time to minimize down time.

Sample Configuration - Two Dedicated Routes

Figure D-2 shows a sample configuration with two dedicated routes, followed by the statements that apply to it.

Figure D-2 Dual IP Configuration Sample - Different Subnetworks



Setup Statements and Displays

The following information shows the statements and displays used in [Figure D-2](#).

HSC LMUPDEF PARMLIB Member

```
OPTION TITLE('DUAL TCPIP CONNECTING TO SL8500')
LMUPATH ACS(00) LMUADDR(129.80.71.83,172.27.2.5)
```

TCP/IP Profile Data Set

```

; OSA CARD #1
DEVICE ECCQD01 MPCIPA NONROUTER AUTORESTART
LINK ZIPBMVS IPAQENET ECCQD01

; OSA CARD #2
DEVICE ECCQA01 MPCIPA NONROUTER AUTORESTART
LINK ZIPB2MVS IPAQENET ECCQA01

HOME
  129.80.16.239 ZIPBMVS
  129.80.65.203 ZIPB2MVS

BEGINROUTES
;
      NETWORK MASK  FIRSTHOP      LINKNAME      PACKETSIZE
ROUTE 129.80.16.0/24      =          ZIPBMVS      MTU 1492
ROUTE 129.80.65.0/24      =          ZIPB2MVS     MTU 1492
ROUTE 172.27.2.5  HOST      129.80.16.254 ZIPBMVS      MTU 1492
ROUTE 129.80.71.83 HOST      129.80.65.254 ZIPB2MVS     MTU 1492
ROUTE DEFAULT 129.80.16.254          ZIPBMVS      MTU 1492
ROUTE DEFAULT 129.80.65.254          ZIPB2MVS     MTU 1492

ENDROUTES

INCLUDE ZIP.TCPIP.PROFILES(COMMON)

START ECCQD01
START ECCQA01

```

TCP/IP Console Displays

```

D TCPIP,TCPIP,NETSTAT,DEV
EZZ2500I NETSTAT CS V1R4 TCPIP 229
DEVNAME: LOOPBACK          DEVTYPE: LOOPBACK
DEVSTATUS: READY
LNKNAME: LOOPBACK          LNKTYPE: LOOPBACK   LNKSTATUS: READY
NETNUM: 0  QUESIZE: 0
BYTESIN: 1781074          BYTESOUT: 1781074
ACTMTU: 65535
BSD ROUTING PARAMETERS:
MTU SIZE: 00000          METRIC: 00
DESTADDR: 0.0.0.0        SUBNETMASK: 0.0.0.0
MULTICAST SPECIFIC:
MULTICAST CAPABILITY: NO

DEVNAME: ECCQD01          DEVTYPE: MPCIPA
DEVSTATUS: READY          CFGROUTER: NON  ACTROUTER: NON
LNKNAME: ZIPBMVS          LNKTYPE: IPAQENET   LNKSTATUS: READY
NETNUM: 0  QUESIZE: 0  SPEED: 0000001000
BYTESIN: 34704496          BYTESOUT: 11207410
IPBROADCASTCAPABILITY: NO
ARPOFFLOAD: YES  ARPOFFLOADINFO: YES
ACTMTU: 8992
INBPERF: BALANCED
BSD ROUTING PARAMETERS:
MTU SIZE: 00000          METRIC: 00
DESTADDR: 0.0.0.0        SUBNETMASK: 255.255.255.0
MULTICAST SPECIFIC:
MULTICAST CAPABILITY: YES
GROUP          REFCNT
-----
224.0.0.1      0000000001

```

```

DEVNAME: ECCQA01          DEVTYPE: MPCIPA
DEVSTATUS: READY         CFGROUTER: NON  ACTROUTER: NON
LNKNAME: ZIPB2MVS        LNKTYPE: IPAQENET  LNKSTATUS: READY
NETNUM: 0  QUESIZE: 0    SPEED: 000000100
BYTESIN: 147508801      BYTESOUT: 2188246
IPBROADCASTCAPABILITY: NO
ARPOFFLOAD: YES  ARPOFFLOADINFO: YES
ACTMTU: 1492
INBPERF: BALANCED
BSD ROUTING PARAMETERS:
MTU SIZE: 00000          METRIC: 00
DESTADDR: 0.0.0.0       SUBNETMASK: 255.255.255.0
MULTICAST SPECIFIC:
MULTICAST CAPABILITY: YES
GROUP          REFCNT
-----
224.0.0.1     0000000001
3 OF 3 RECORDS DISPLAYED
D TCPIP,TCPIP,NETSTAT,ROUTE
EZZ2500I NETSTAT CS V1R4 TCPIP 250
DESTINATION      GATEWAY          FLAGS      REFCNT  INTERFACE
DEFAULT          129.80.16.254    UGS        000003  ZIPBMVS
DEFAULT          129.80.65.254    UGS        000002  ZIPB2MVS
127.0.0.1        0.0.0.0          UH         000003  LOOPBACK
129.80.16.0      0.0.0.0          US         000002  ZIPBMVS
129.80.16.239   0.0.0.0          UH         000000  ZIPBMVS
129.80.65.0     0.0.0.0          US         000000  ZIPB2MVS
129.80.65.203   0.0.0.0          UH         000000  ZIPB2MVS
129.80.71.83    129.80.65.254    UGHS       000001  ZIPB2MVS
172.27.2.5      129.80.16.254    UGHS       000001  ZIPBMVS
9 OF 9 RECORDS DISPLAYED

```

HSC Display Showing SL8500 Connection

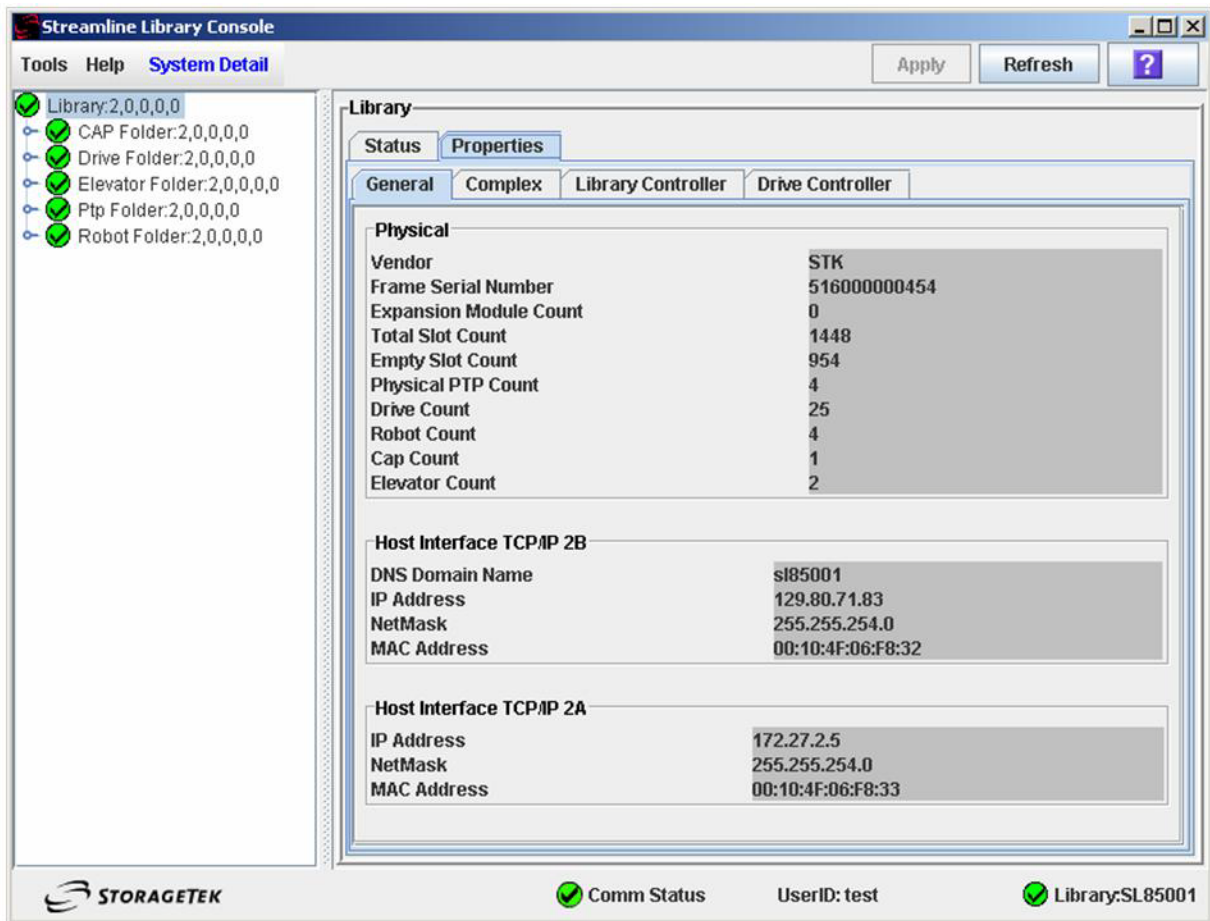
```

D ACS
SLS0000I D ACS
SLS1000I ACS 00 STATUS: CONNECTED 334
ACTIVE QUEUE ELEMENTS          2
COMPATIBILITY LEVELS: HSC=13, LMU=13
DUAL LMU NOT CONFIGURED
IP ADDR 129.80.71.83 ONLINE
IP ADDR 172.27.2.5 ONLINE
SCRATCH VOLUMES AVAILABLE..... 210
FREE CELLS AVAILABLE..... 2008
D TCPIP,TCPIP,NETSTAT,HOME
EZZ2500I NETSTAT CS V1R4 TCPIP 252
HOME ADDRESS LIST:
ADDRESS      LINK          FLG
129.80.16.239  ZIPBMVS      P
129.80.65.203  ZIPB2MVS
127.0.0.1     LOOPBACK
3 OF 3 RECORDS DISPLAYED

```

Figure D-3 shows an SLConsole display of ports 2A and 2B, including IP addresses.

Figure D-3 SLConsole Display for SL8500 Ports 2A and 2B

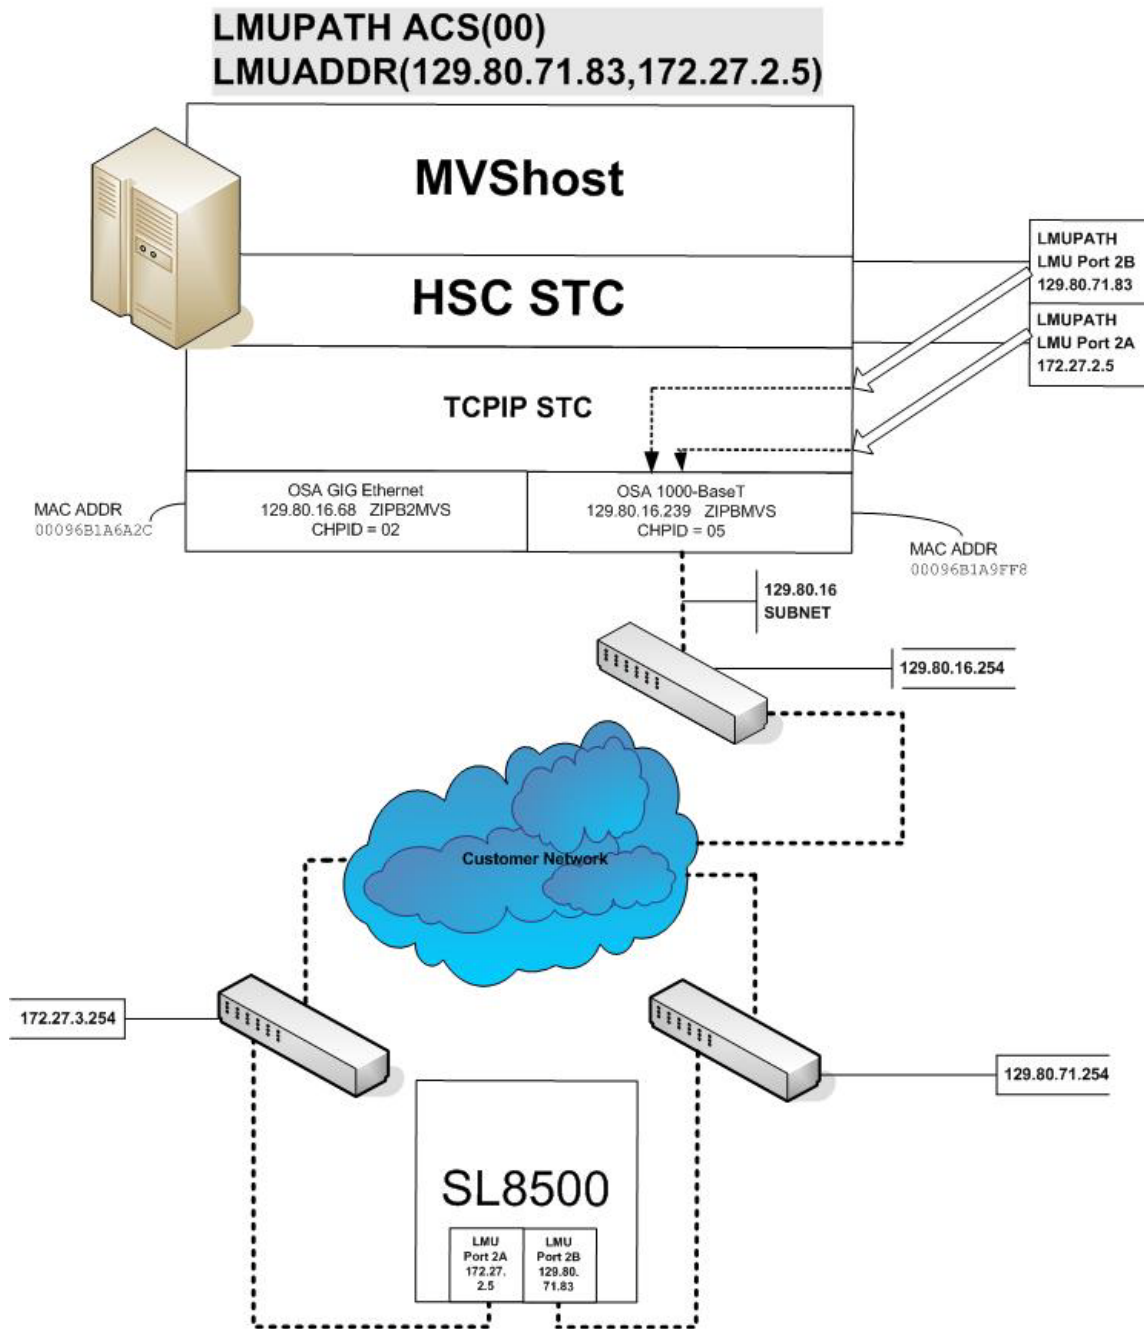


Two SL8500 Network Connections - One Mainframe IP Address

To establish one host IP route from the HSC to two routes to the SL8500, use the process described in this section. To configure the SL8500 dual TCP/IP feature, refer to the *StreamLine SL8500 Modular Library System Dual TCP/IP Feature* document.

Figure D-4 shows dual IP (one host IP) with two SL8500 network connections.

Figure D-4 Dual IP (one host IP) with Two SL8500 Network Connections



1. Complete the Network Entries Worksheet (see [Table D-2](#) for a sample) for each set of routes to the SL8500. The worksheet can be found in the StreamLine SL8500 Modular Library System Dual TCP/IP Feature document.

Table D-2 Network Entries Worksheet

Port Number	Description	IP Address
2A	HBC Card Port 2A	172.27.2.5
2A	Gateway Port 2A	172.27.3.254

Table D-2 (Cont.) Network Entries Worksheet

Port Number	Description	IP Address
2A	Netmask	/23
2A	Host to Port 2A	129.80.16.239
2B	HBC Car Port 2B	129.80.71.83
2B	Gateway Port 2B	129.80.71.254
2B	Netmask	/23
2B	Host to Port 2B	129.80.65.203

- Define two dedicated static routes to the SL8500 destination port (2A and 2B) IP addresses over one router.

Example:

```
BEGINROUTES
;   DESTINATION                FIRSTHOP      LINKNAME      PACKETSIZE
ROUTE 129.80.16.0/24          =             &SYSNAME.MVS  MTU 1492
ROUTE 172.27.2.5 HOST        129.80.16.254 &SYSNAME.MVS  MTU 1492
ROUTE 129.80.71.83 HOST      129.80.16.254 &SYSNAME.MVS  MTU 1492
ROUTE DEFAULT                 129.80.16.254 &SYSNAME.MVS  MTU 1492

ENDROUTES
```

- Define a second LMUADDR parameter for the port 2A IP address on the SL8500.

Example:

```
LMUPATH ACS(00) LMUADDR(129.80.71.83,172.27.2.5)
```

- Enter the LMUPDEF command containing the LMUPATH statements that define the host name or IP address for each ACS.

```
LMUPDEF DSN('xxx.xxx.xxx(xxx)')
```

In the following example, LMUPDEF loads LMUPATH parameters from YOUR.DSN(MEMBER).

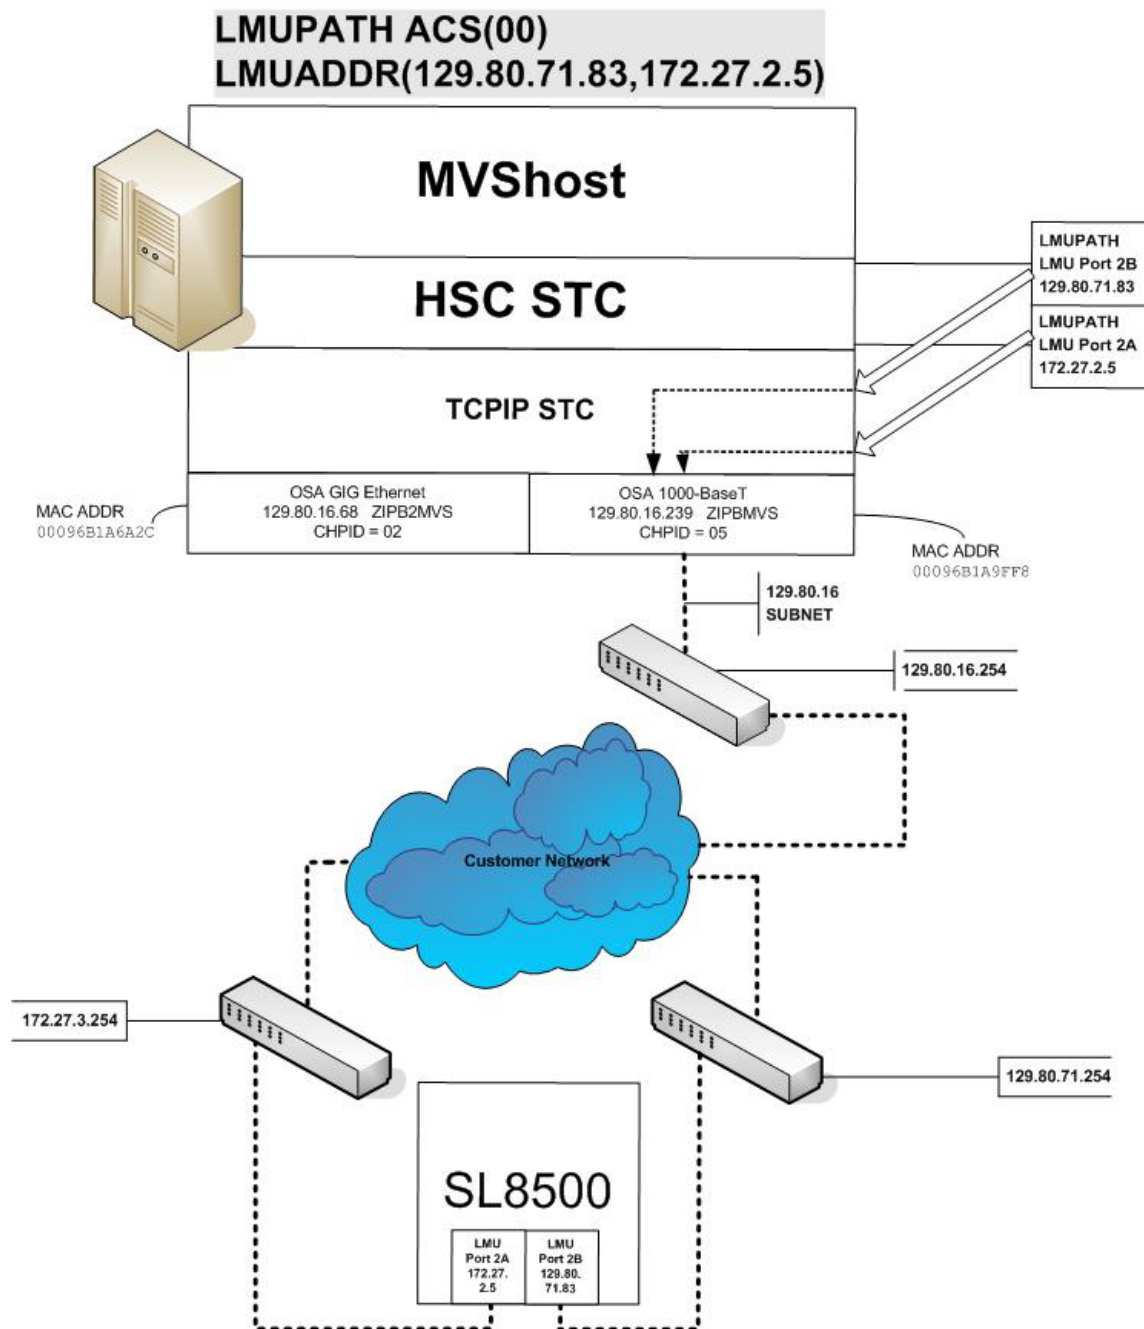
```
LMUPDEF DSN('YOUR.DSN(MEMBER)')
```

- Allow the trained SL8500 service representative to enter the network connections to the SL8500 library for either port 2A and 2B, whichever is applicable.
- Vary the ACS offline and back online to pick up the revised LMUPATH statement that includes the second connection. This can be done one host at a time to minimize down time.

Sample Configuration - One Host IP, Two SL8500 Network Connections

Figure D-5 shows a sample dual IP (one host IP) configuration with two SL8500 network connections, followed by the statements that apply to it.

Figure D-5 Dual IP (one host IP) with Two SL8500 Network Connections



Setup Statements and Displays

The following information shows the statements and displays used in [Figure D-5](#).

HSC LMUPDEF PARMLIB Member

```
OPTION TITLE('DUAL TCPIP CONNECTING TO SL85001')
LMUPATH ACS(00) LMUADDR(129.80.71.83,172.27.2.5)
```

TCP/IP Profile Data Set

```
; OSA CARD #1
DEVICE ECCQD01 MPCIPA NONROUTER AUTORESTART
```

```

LINK    ZIPBMVS IPAQENET ECCQD01

HOME
  129.80.&IPADDR1 &SYSNAME.MVS

BEGINROUTES
;
      NETWORK MASK  FIRSTHOP      LINKNAME      PACKETSIZE
ROUTE 129.80.16.0/24      =          &SYSNAME.MVS  MTU 1492
ROUTE 172.27.2.5 HOST    129.80.16.254 &SYSNAME.MVS  MTU 1492
ROUTE 129.80.71.83 HOST  129.80.16.254 &SYSNAME.MVS  MTU 1492
ROUTE DEFAULT

ENDROUTES

INCLUDE ZIP.TCPIP.PROFILES(COMMON)

START ECCQD01

```

TCP/IP Console Displays

```

D TCPIP, ,N,DEV
EZZ2500I NETSTAT CS V1R4 TCPIP 931
DEVNAME: LOOPBACK          DEVTYPE: LOOPBACK
DEVSTATUS: READY
LNKNAME: LOOPBACK          LNKTYPE: LOOPBACK    LNKSTATUS: READY
NETNUM: 0  QUESIZE: 0
BYTESIN: 2136824          BYTESOUT: 2136824
ACTMTU: 65535
BSD ROUTING PARAMETERS:
  MTU SIZE: 00000          METRIC: 00
  DESTADDR: 0.0.0.0        SUBNETMASK: 0.0.0.0
MULTICAST SPECIFIC:
  MULTICAST CAPABILITY: NO
DEVNAME: ECCQD01          DEVTYPE: MPCIPA
DEVSTATUS: READY          CFGROUTER: NON  ACTROUTER: NON
LNKNAME: ZIPBMVS          LNKTYPE: IPAQENET  LNKSTATUS: READY
BYTESIN: 48605838        BYTESOUT: 9790950
NETNUM: 0  QUESIZE: 0  SPEED: 0000001000
IPBROADCASTCAPABILITY: NO
ARPOFFLOAD: YES  ARPOFFLOADINFO: YES
ACTMTU: 8992
INBPERF: BALANCED
BSD ROUTING PARAMETERS:
  MTU SIZE: 00000          METRIC: 00
  DESTADDR: 0.0.0.0        SUBNETMASK: 255.255.255.0
MULTICAST SPECIFIC:
  MULTICAST CAPABILITY: YES
  GROUP                    REFCNT
  -----
2 OF 2 RECORDS DISPLAYED

```

```

D TCPIP, ,N,ROUTE
EZZ2500I NETSTAT CS V1R4 TCPIP 933
DESTINATION      GATEWAY      FLAGS      REFCNT  INTERFACE
DEFAULT          129.80.16.254  UGS        000001  ZIPBMVS
127.0.0.1        0.0.0.0      UH         000004  LOOPBACK
129.80.16.0      0.0.0.0      US         000000  ZIPBMVS
129.80.16.239    0.0.0.0      UH         000000  ZIPBMVS
129.80.71.83     129.80.16.254 UGHS       000001  ZIPBMVS
172.27.2.5       129.80.16.254 UGHS       000001  ZIPBMVS
6 OF 6 RECORDS DISPLAYED

```

```

D TCPIP,,N,HOME
EZZ2500I NETSTAT CS V1R4 TCPIP 935
HOME ADDRESS LIST:
ADDRESS          LINK          FLG
129.80.16.239   ZIPBMVS      P
127.0.0.1      LOOPBACK
2 OF 2 RECORDS DISPLAYED
    
```

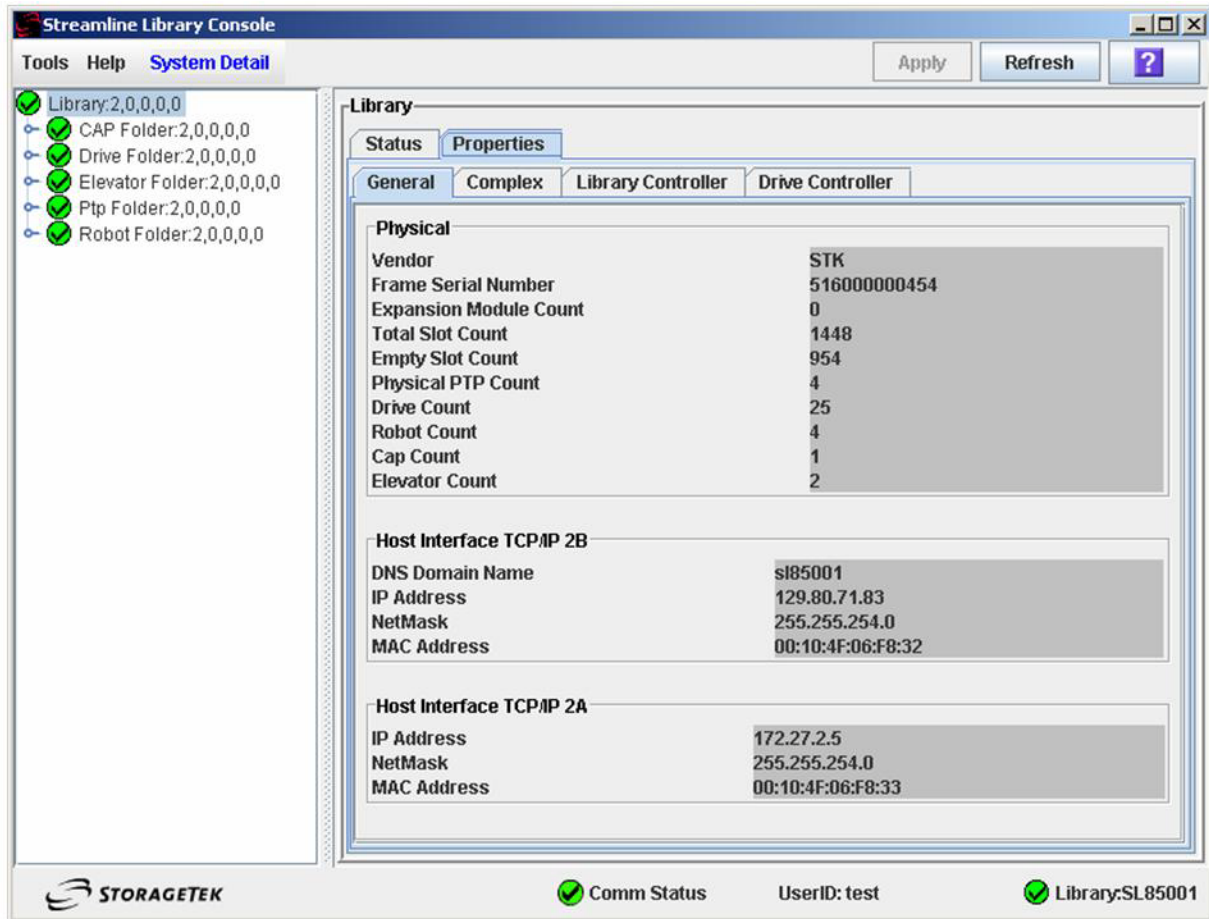
HSC Display Showing SL8500 Connection

```

D ACS
SLS0000I D ACS
SLS1000I ACS 00 STATUS: CONNECTED 942
ACTIVE QUEUE ELEMENTS          1
COMPATIBILITY LEVELS: HSC=13, LMU=13
DUAL LMU NOT CONFIGURED
IP ADDR 129.80.71.83 ONLINE
IP ADDR 172.27.2.5 ONLINE
SCRATCH VOLUMES AVAILABLE..... 210
FREE CELLS AVAILABLE.....      2007
    
```

Figure D-6 shows an SLConsole display of ports 2A and 2B, including IP addresses.

Figure D-6 SLConsole Display for SL8500 Ports 2A and 2B



Multiple SL8500 Library Connections

This section discusses the following:

- ["Connecting to Multiple SL8500 Libraries in an ACS"](#)
- ["Sample Configuration - Four SL8500 Network Connections, Four Mainframe IP Addresses"](#)

Connecting to Multiple SL8500 Libraries in an ACS

When SL8500 3.9x or higher firmware is installed, the HSC can connect to more than one SL8500 in an ACS (library complex).

The HSC supports up to four connections to an ACS. Some possible connections are:

- four connections to four SL8500s
- two connections to each of two SL8500s
- two connections to one SL8500 and two connections to two other SL8500s.

When the HSC establishes two connections to one SL8500, the HSC provides dual TCP/IP or multiple TCP/IP connectivity. Refer to the ["TCP/IP Communications - Important Considerations"](#) discussion for more information.

Note:

- To optimize library performance and minimize inter-library communication among SL8500s, connect to the libraries with the most activity. The HSC distributes communications to the libraries evenly.
 - LMUPATH control statements are used to define network LMU attachments.
-
-

The following example shows multiple LMUADDR parameters consisting of four IP addresses. In this case, the first, second, third, and fourth IP addresses indicate a TCP/IP connection to each of four separate SL8500 libraries connected in ACS 00:

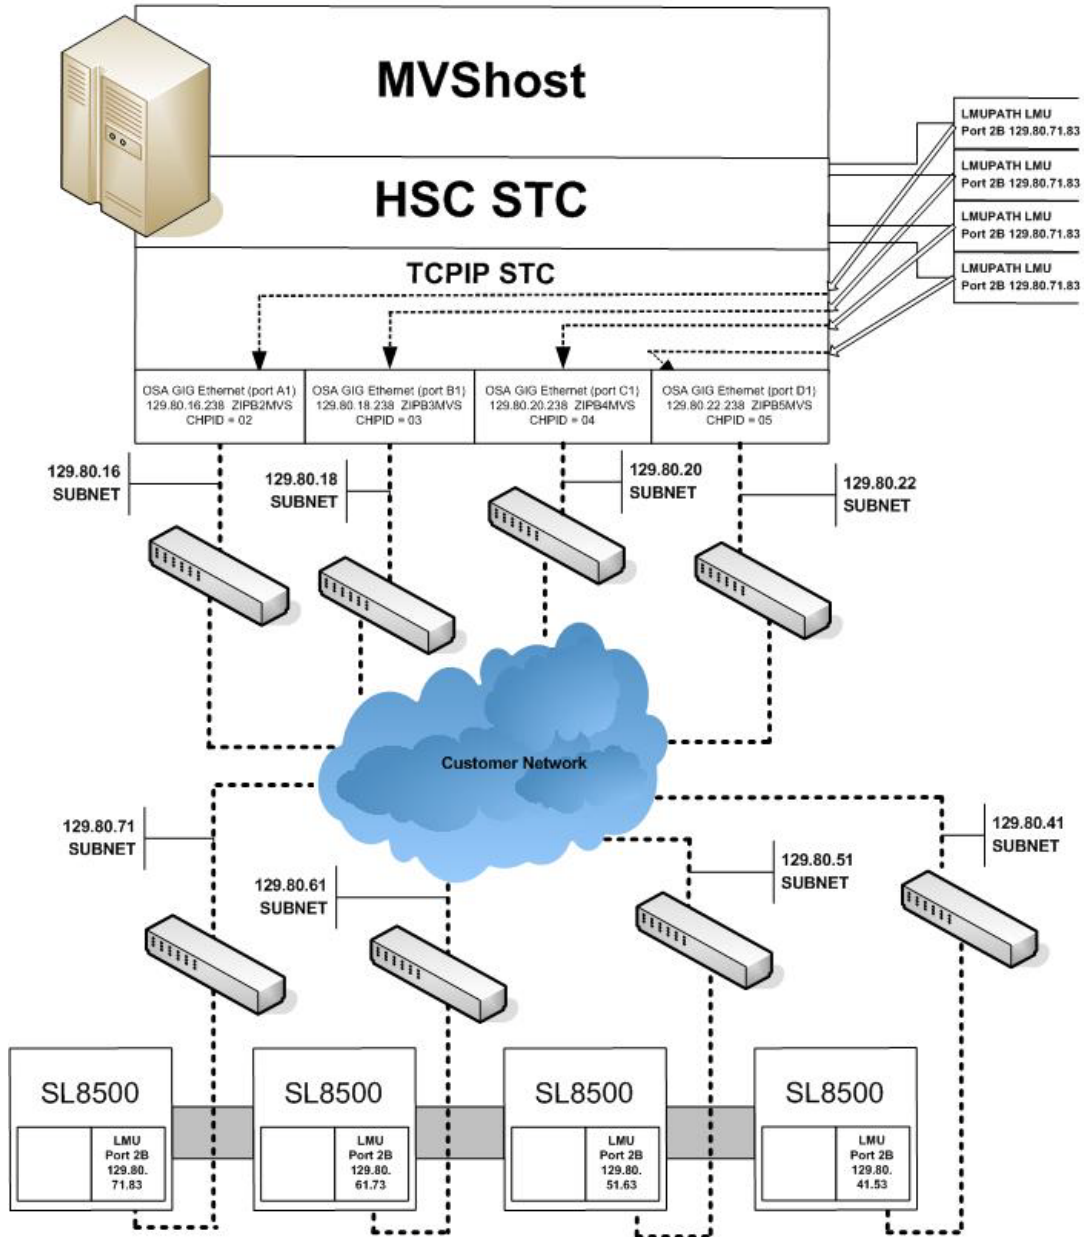
```
LMUPATH ACS(00)+
LMUADDR(123.456.789.012,123.456.789.013,123.456.789.014,123.456.789.015)
```

Sample Configuration - Four SL8500 Network Connections, Four Mainframe IP Addresses

[Figure D-7](#) shows a sample of a four IP (four host IPs) configuration with four SL8500 network connections. This configuration applies to SL8500 3.x or higher firmware.

Figure D-7 Four IP Addresses (four host IPs) with Four SL8500 Network Connections

LMUPATH ACS(00)
LMUADDR(129.80.71.83,129.80.61.73,129.80.51.63,129.80.41.53)



Powering Down the SL8500 - HSC Requirements

Prior to powering down the SL8500, the HSC requires that the following procedure must be followed:

1. Vary the drives offline at the MVS host level.
2. Modify the LSM(s) offline (see the MODify Command in the *ELS Command, Control Statement, and Utility Reference*).

For specific information about powering off the SL8500, refer to the *SL8500 User's Guide*.

Multiple TCP/IP Redundant Electronics

The LMUPATH control statement allows users to define network Library Controller attachments. In a Multiple TCP/IP redundant electronics environment for an SL8500 complex, specify up to 32 network connections using the LMUADDR parameter. When specified, HSC automatically adjusts when Redundant Electronics is installed.

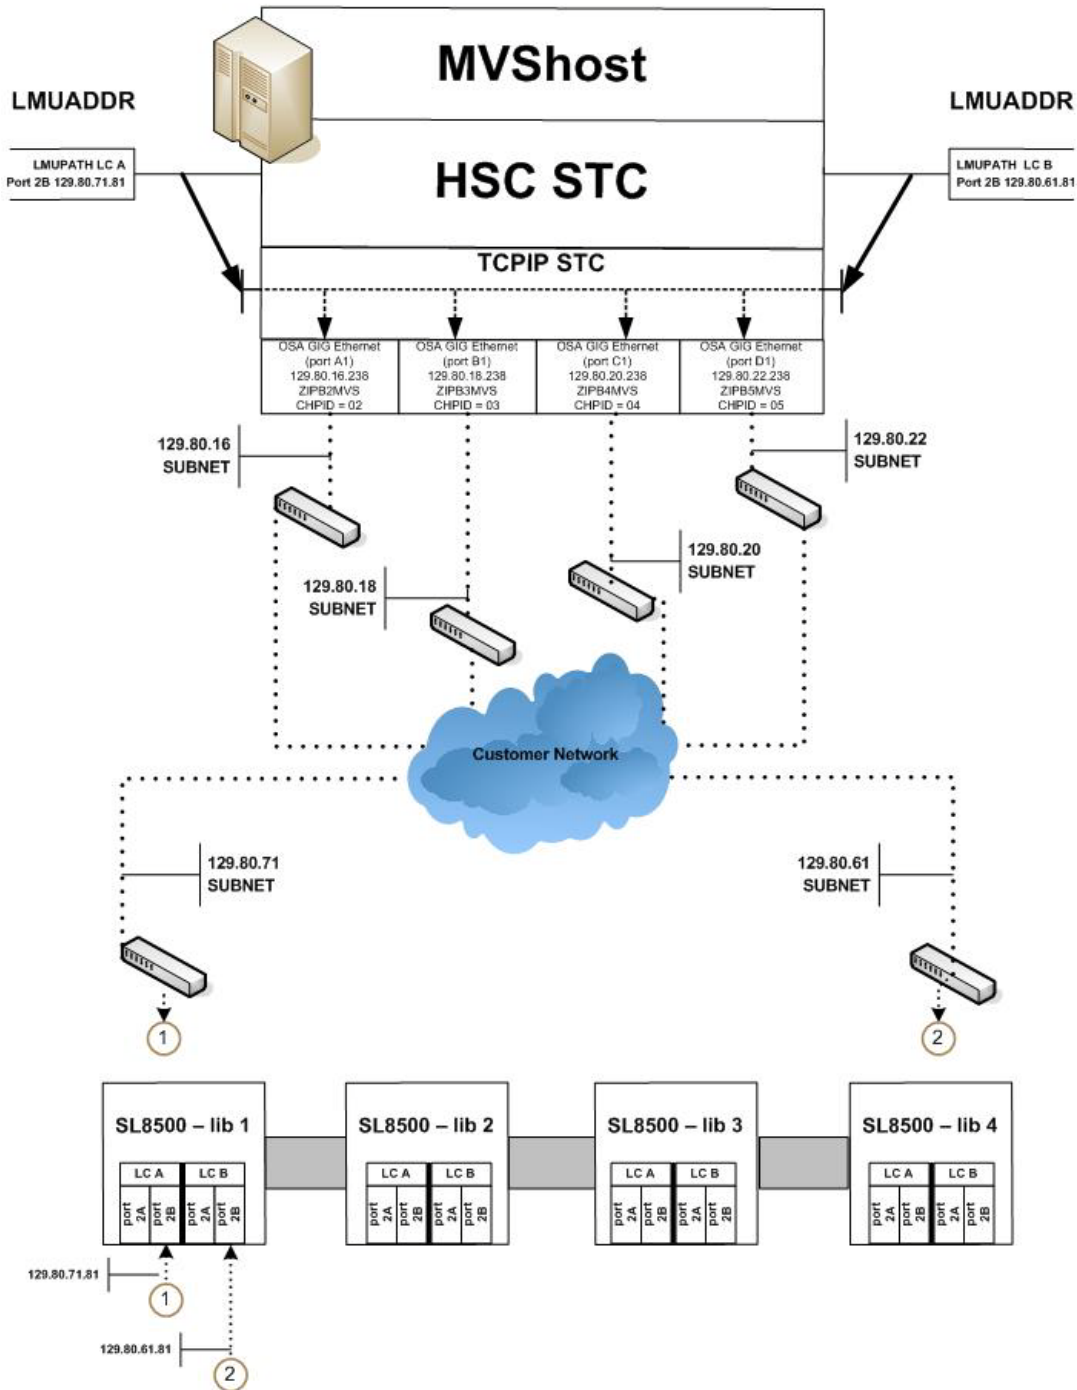
Note: Vary the ACS offline, updating the LMUPDEF file to accommodate redundant electronics, read the new LMUPDEF file, and vary the ACS back online to pick up the revised LMUPATH statement that includes the added connections.

Sample Configuration #1 - One SL8500 Network Single Redundant Pair Connections, Four Mainframe IP Addresses

[Figure D-8](#) shows a sample four IP (four host IPs) configuration with one single redundant pair SL8500 network connections. This configuration applies to SL8500 6.0 or higher firmware.

Figure D-8 Four IP (four host IPs) with One Single Redundant Pair SL8500 Network Connections

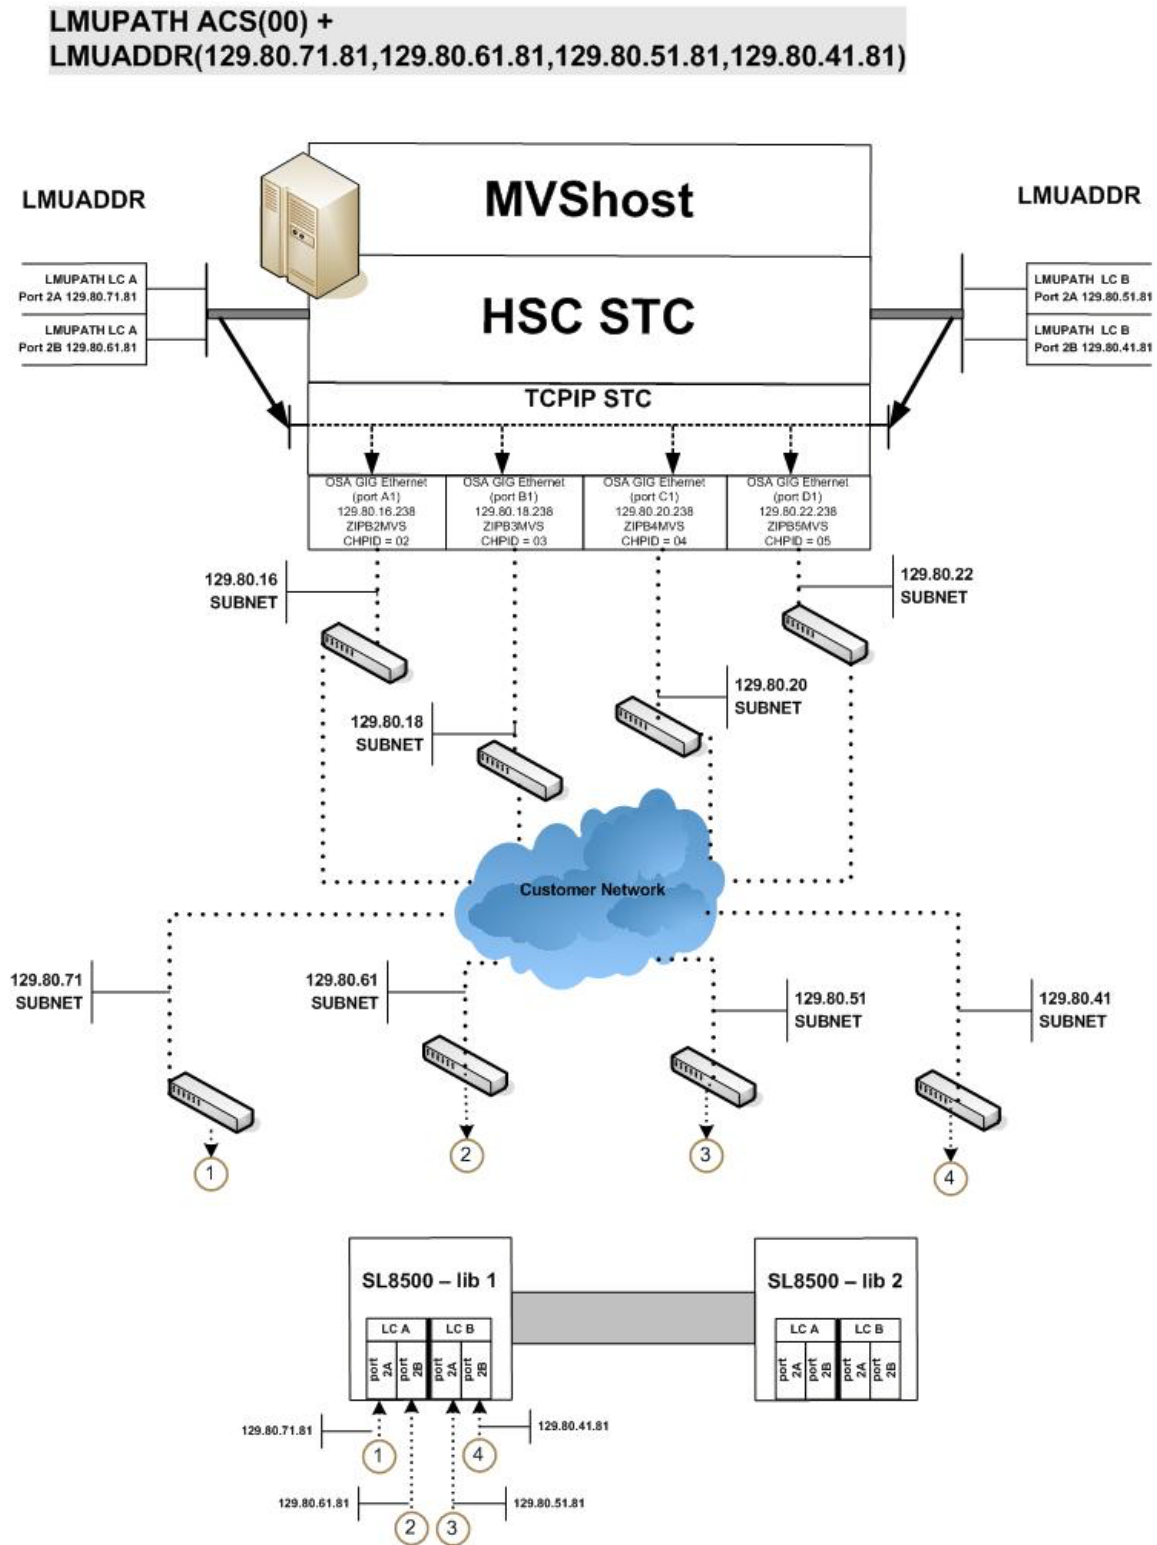
LMUPATH ACS(00) +
LMUADDR(129.80.71.81,129.80.61.81)



Sample Configuration #2 - One SL8500 Network Dual Redundant Pair Connections, Four Mainframe IP Addresses

Figure D-9 shows a sample four IP (four host IPs) configuration with one dual redundant pair SL8500 network connections. This configuration applies to SL8500 6.0 or higher firmware.

Figure D-9 Four IP (four host IPs) with One Dual Redundant Pair SL8500 Network Connections

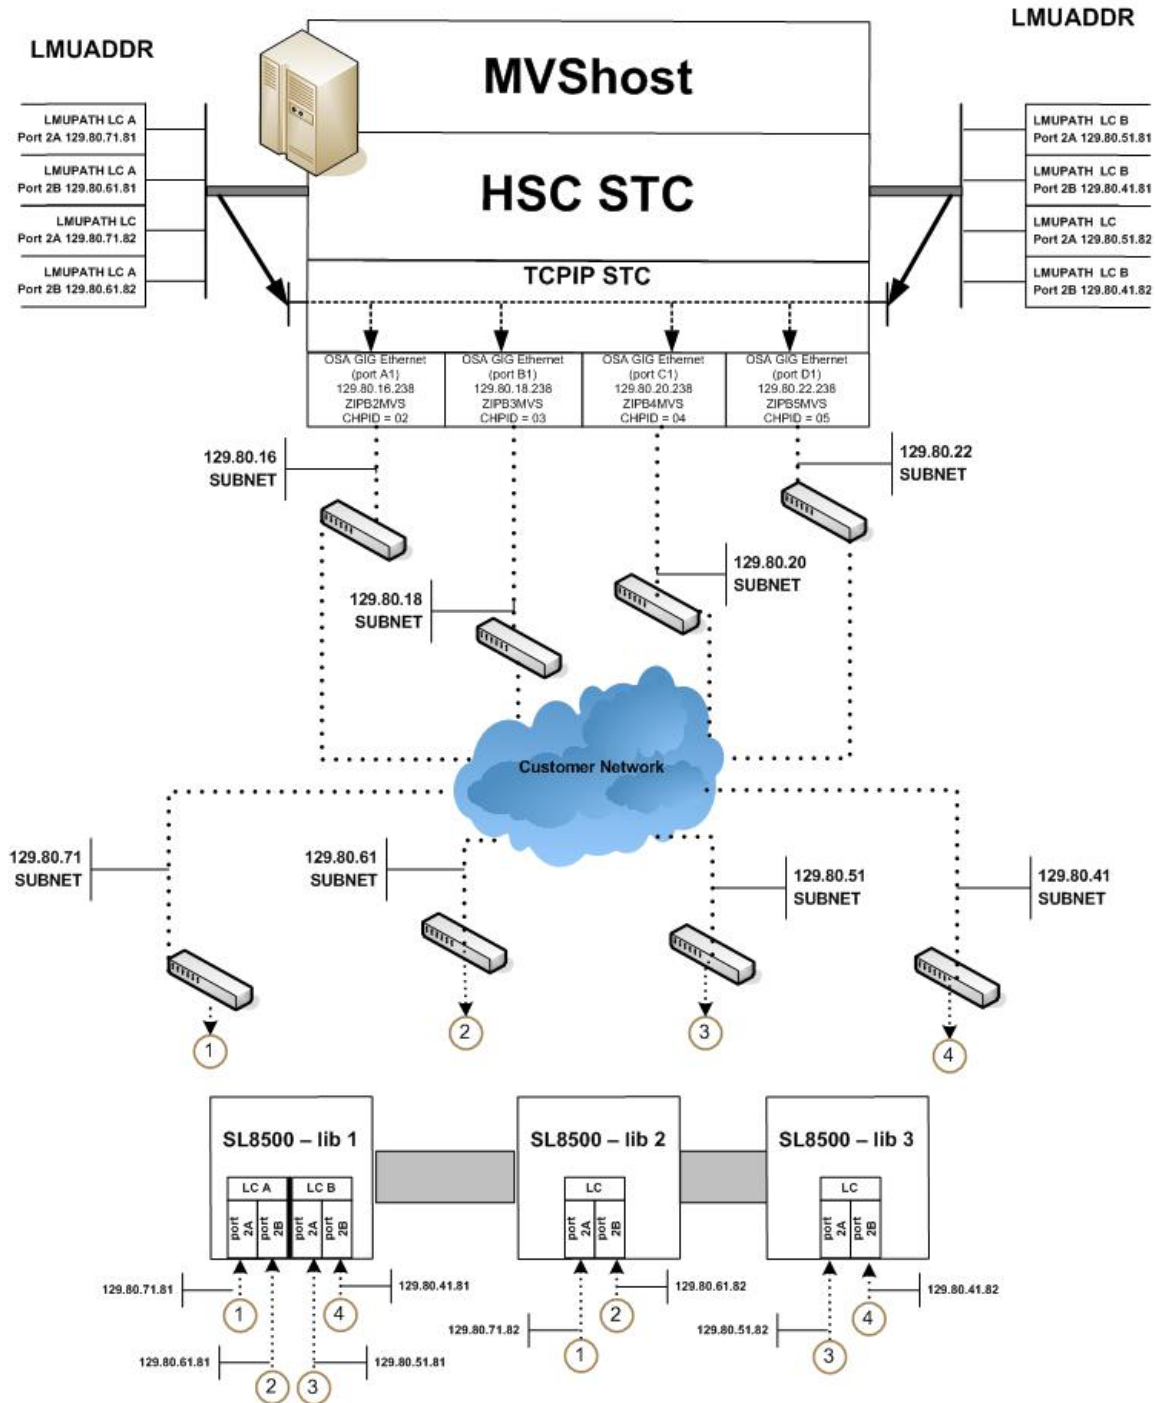


Sample Configuration #3 - One SL8500 Network Dual Redundant Pair and 2 Dual TCP/IP Connections, Four Mainframe IP Addresses

[Figure D-10](#) shows a sample four IP (four host IPs) configuration with one dual redundant pair SL8500 and two dual TCP/IP network connections. This configuration applies to SL8500 6.0 or higher firmware.

Figure D-10 Four IP (four host IPs) with One Dual Redundant Pair SL8500 and 2 Dual TCP/IP Network Connections

**LMUPATH ACS(00) +
 LMUADDR(129.80.71.81,129.80.61.81,129.80.51.81,129.80.41.81, +
 129.80.71.82,129.80.61.82,129.80.51.82,129.80.41.82)**



Troubleshooting Library Connectivity Problems

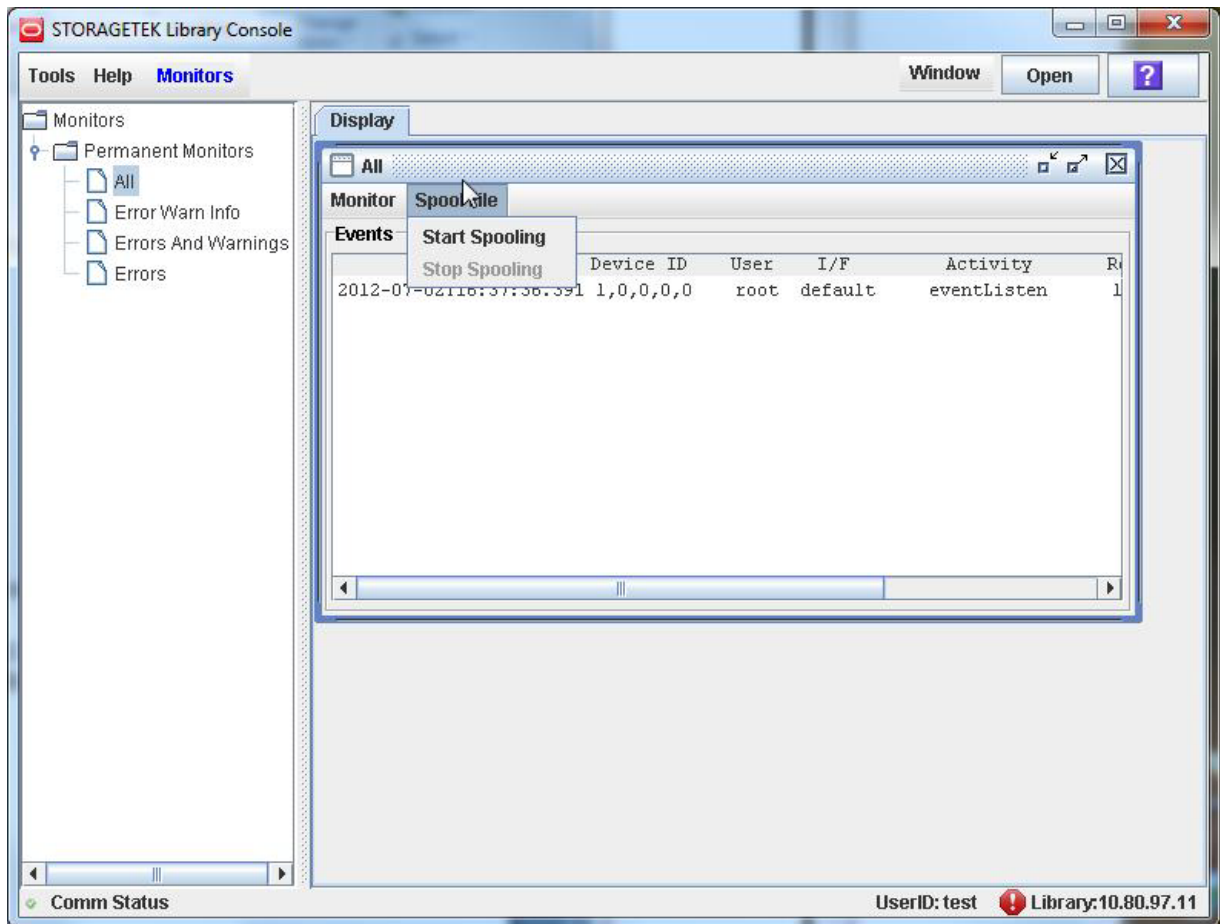
This section covers procedures on helping you to provide important data to Oracle support when Library connectivity problems occur.

Basic Tracing

The following procedure tells how to produce diagnostic data for basic SL3000/SL8500 host connectivity problems.

1. Using the SLConsole, start a monitor trace (select All) and start spooling to an output file as shown in the example in [Figure D-11](#).

Figure D-11 Example SLC Console (Monitor Trace)



2. Update the HSC PARMLIB member (in bold) to add LMU tracing on as shown in the example.

```
CDSDEF DSN1 (HSC1 . PRIM) , DSN2 (HSC1 . BKUP) , DSN3 (HSC1 . STDBY)
MGMTDEF DSN (HSC . PARMS) COMP METH VTAMFEAT VSM (ADVMGMT)
CAPP 9,01:00:00,AUTO
MNTD AUTOCLN (ON) EJCTAUTO (ON)
/* turn LMU trace on */
TR LMU
```

3. Start GTF to trace the HSC job as shown in the JCL example below.

```
//GTFTEST PROC MEMBER=GTFTEST
//IEFPROC EXEC PGM=AHLGTF, PARM='MODE=EXT, DEBUG=NO, TIME=YES',
//          REGION=2880K, ACCT=505341SPKC
//IEFRDER DD  DSN=NAME=GTF.OUTPUT, DISP=SHR /* output trace file */
//SYSLIB  DD  DSN=TEST.PARMLIB(&MEMBER), DISP=SHR
```

Or, use a GTF SYSLIB member as shown in the following example.

```
TRACE=USR, JOBNAMEP
JOBNAME=(HSCT*) /* HSCT* = HSC jobname prefix (jobname example: HSCT1234) */
END
```

4. After a problem occurs:
 - a. Stop GTF.
 - b. Stop the spooling in the monitor trace.
 - c. Stop HSC.
5. Send to the monitor trace, GTF output files, and HSC log to StorageTek/Oracle support.

Packet Tracing

A packet trace, which traces the actual data packets that are being transferred across the line, can be helpful when trying to determine if data is actually being sent across properly. The following procedure tells how to run a packet trace.

1. Create an external writer proc in SYS1.PROCLIB(CTTCP) as shown in the following example.

```
//CTTCP PROC
//CTTCP EXEC PGM=ITTTTCWR
//TRCOUT01 DD DSN=Your_data_set, UNIT=SYSDA,
//          SPACE=(CYL,(100),,CONTIG), DISP=(NEW,CATLG)
```

2. Enter the following MVS commands:

```
TRACE CT, WTRSTART=CTTCP
TRACE CT, ON, COMP=SYSTCPDA, SUB=(your TCP/IP started task name)
R n, WTR=CTTCP, END
D TRACE, COMP=SYSTCPDA, SUB=(your TCP/IP started task name)
```

Note: This will verify that trace was started successfully.

```
V TCPIP, your TCP/IP started task name, PKT, ON, IP=xx.xx.xx.xx
```

Where *xx.xx.xx.xx* is the IP address you want to trace.

3. Run the trace until the error occurs.
4. Enter the following MVS commands:

```
V TCPIP, your TCP/IP started task name, PKT, OFF
TRACE CT, ON, COMP=SYSTCPDA, SUB=(your TCP/IP started task name)
R n, WTR=DISCONNECT, END
TRACE CT, OFF, COMP=SYSTCPDA, SUB=(your TCP/IP started task name)
TRACE CT, WTRSTOP=CTTCP, FLUSH
```

5. To format the trace, use IPCS and the following command:

```
CTRACE COMP (SYSTCPDA) FULL
```

Or use IPCS menu option 2, Option 7, Option 1, Option D and then enter component name SYSTCPDA and report type FULL.

Configuring SL3000 Library Communications

The SL3000 is different from other LSMs, and the way it communicates with HSC is also different, as explained in this appendix.

TCP/IP Communications - Important Considerations

There are two types of TCP/IP connections for SL3000 discussed in this appendix:

- "Shared Networks"
- "SL3000 Dual IP Connections" - A connecting redundancy feature to allow up to two connections to the library.

Note: It is crucial that you **read completely** this section to understand these procedures.

Shared Networks

The following are some examples of issues that can arise when you connect a Streamline library to a shared network.

- A TCP/IP-connected library can handle standard host traffic but it cannot resolve floods of Address Resolution Protocol (ARP) broadcasts. **For this reason, it is best to attach the library on a controlled network, such as behind a switch or router.**

Later generation networks, such as 1000Base-T and Gig-E, support earlier communication modes; however, devices that are communicating with the library may transmit data at bandwidths that could overwhelm the library.

It is best to attach the library on a controlled network, such as with a switch that can isolate the library from network broadcasts.

- When you connect the library on shared networks and broadcasts are sent to all network nodes, they may also be directed to the library (even though it does not need them).

During the time the library is receiving these irrelevant broadcasts, it cannot receive from or respond to other requests in a timely fashion. This heavy broadcast traffic on the network can saturate the library to the point that, to the host, it may appear that the TCP/IP connection has been lost. It may also result in Overdue Response Handler (ORH) messages on the host system.

- Heavy network traffic can also overwhelm the Ethernet controller causing the processor to continuously reset and re-initialize the controller, then recover the host-to-library communications.

SL3000 Dual IP Connections

The LMUPATH control statement allows users to define network LMU attachments. In a dual IP connection environment for an SL3000, specify a second LMUADDR parameter to define dual IP. The HSC automatically determines whether or not the connection is dual IP or dual LMU.

Note: Vary the ACS offline and back online to pick up the revised LMUPATH statement that includes the second connection.

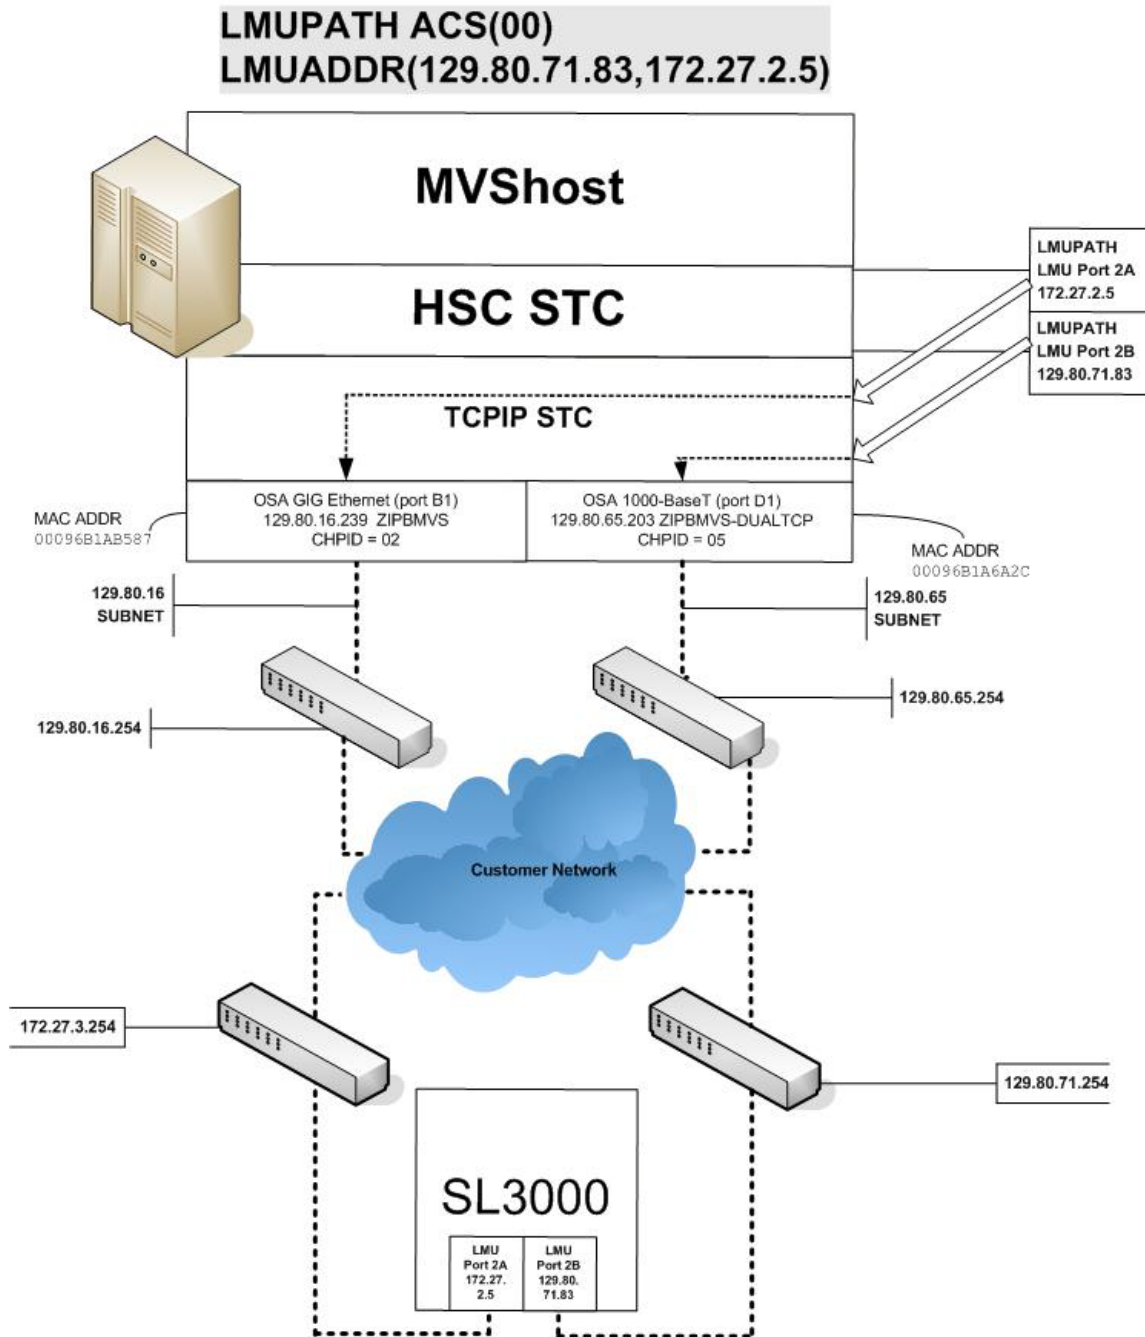
[Figure E-1](#) and the example that follows it shows a dual IP connection for HSC.

Two SL3000 Network Connections - Two Mainframe IP Addresses

To establish two dedicated routes on different subnetworks from the HSC to two dedicated routes to the SL3000, use the process described in this section. To configure the SL3000 dual TCP/IP feature, refer to the *StreamLine SL3000 Modular Library System Installation Guide*.

[Figure E-1](#) shows dual IP connections with two SL3000 network connections.

Figure E-1 Dual IP with Two SL3000 Network Connections - Different Subnetworks



1. Create and complete the Network Entries Worksheet (see Table E-1 for a sample) for each set of two dedicated routes to the SL3000. The worksheet can be found in the *SL3000 Modular Library System Installation Guide*.

Table E-1 Network Entries Worksheet

Port Number	Description	IP Address
2A	HBC Card Port 2A	172.27.2.5
2A	Gateway Port 2A	172.27.3.254

Table E-1 (Cont.) Network Entries Worksheet

Port Number	Description	IP Address
2A	Netmask	/23
2A	Host to Port 2A	129.80.16.239
2B	HBC Card Port 2B	129.80.71.83
2B	Gateway Port 2B	129.80.71.254
2B	Netmask	/23
2B	Host to Port 2B	129.80.65.203

- Define a second DEVICE and LINK statement in your TCP/IP profile data set for a second mainframe network connection.

Example:

```
; OSA CARD #1
DEVICE ECCQD01 MPCIPA NONROUTER AUTORESTART
LINK ZIPBMVS IPAQENET ECCQD01
```

```
; OSA CARD #2
DEVICE ECCQA01 MPCIPA NONROUTER AUTORESTART
LINK ZIPB2MVS IPAQENET ECCQA01
```

- Define a second home address in your TCP/IP profile data set.

Example:

```
HOME
129.80.16.239 ZIPBMVS
129.80.65.203 ZIPB2MVS
```

- Define a second router on the second subnetwork in the routing paragraph of your TCP/IP profile data set.

Remember that you must also configure the SL3000 routing tables according to the instructions described in the *SL3000 Modular Library System Installation Guide*.

Example:

```
BEGINROUTES
; NETWORK MASK FIRSTHOP LINKNAME PACKETSIZE
ROUTE 129.80.16.0/24 = ZIPBMVS MTU 1492
ROUTE 129.80.65.0/24 = ZIPB2MVS MTU 1492
ROUTE 172.27.2.5 HOST 129.80.16.254 ZIPBMVS MTU 1492
ROUTE 129.80.71.83 HOST 129.80.65.254 ZIPB2MVS MTU 1492
ROUTE DEFAULT 129.80.16.254 ZIPBMVS MTU 1492
ROUTE DEFAULT 129.80.65.254 ZIPB2MVS MTU 1492
ENDROUTES
```

- Optionally, define two dedicated static routes to the SL3000 destination port (2A and 2B) IP addresses over two different routers.

Example:

```
; NETWORK MASK ROUTER LINKNAME PACKETSIZE
ROUTE S18500-port-2A-IP-Address HOST 129.80.16.254 MVSHOST1 MTU 1492
ROUTE S18500-port-2B-IP-Address HOST 129.80.64.254 MVSHOST2 MTU 1492
```

```

BEGINROUTES
;
NETWORK MASK    FIRSTHOP    LINKNAME    PACKETSIZE
ROUTE 129.80.16.0/24    =    ZIPBMVS    MTU 1492
ROUTE 129.80.65.0/24    =    ZIPB2MVS    MTU 1492
ROUTE 172.27.2.5    HOST    129.80.16.254    ZIPBMVS    MTU 1492
ROUTE 129.80.71.83    HOST    129.80.65.254    ZIPB2MVS    MTU 1492
ROUTE DEFAULT 129.80.16.254    ZIPBMVS    MTU 1492
ROUTE DEFAULT 129.80.65.254    ZIPB2MVS    MTU 1492
ENDROUTES

```

6. Start the second mainframe network connection device.

```
V TCPIP, tcp-stc-name, START, device_name
```

7. Define a second LMUADDR parameter for the port 2A IP address on the SL3000.

Example:

```
LMUPATH ACS(00) LMUADDR(129.80.71.83,172.27.2.5)
```

8. Enter the LMUPDEF command containing the LMUPATH statements that define the host name or IP address for each ACS.

```
LMUPDEF DSN('xxx.xxx.xxx(xxx)')
```

In the following example, LMUPDEF loads LMUPATH parameters from YOUR.DSN(MEMBER).

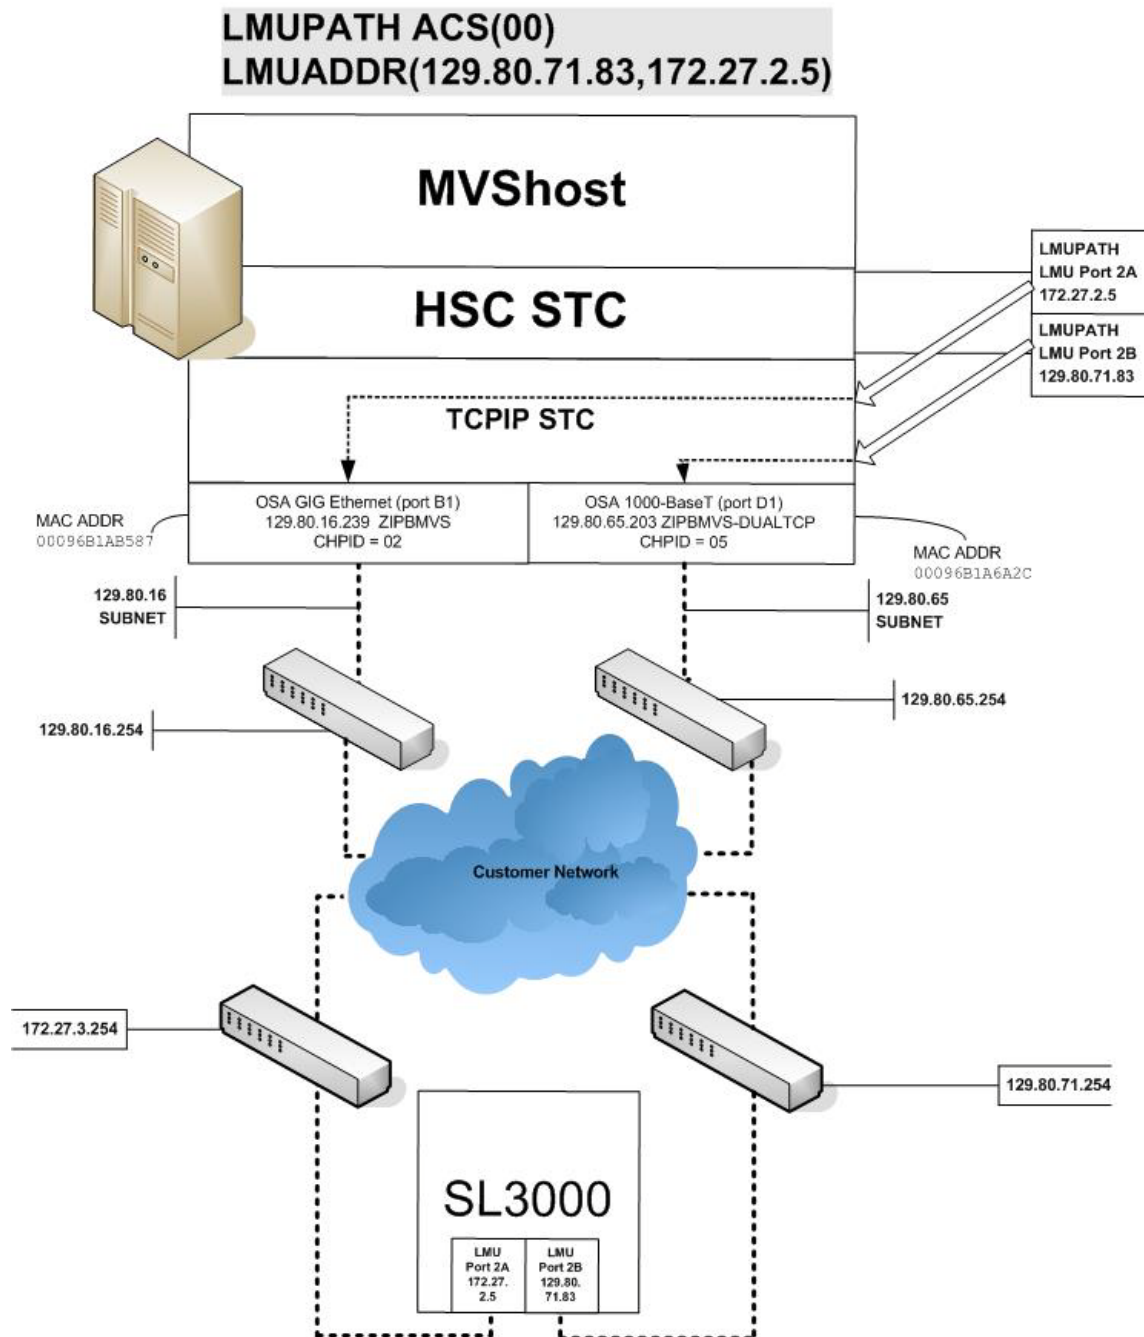
```
LMUPDEF DSN('YOUR.DSN(MEMBER)')
```

9. Allow the trained SL3000 service representative to enter the network connections to the SL3000 library for either port 2A and 2B, whichever is applicable.
10. Vary the ACS offline and back online to pick up the revised LMUPATH statement that includes the second connection. This can be done one host at a time to minimize down time.

Sample Configuration - Two Dedicated Routes

Figure E-2 shows a sample configuration with two dedicated routes, followed by the statements that apply to it.

Figure E-2 Dual IP Configuration Sample - Different Subnetworks



Setup Statements and Displays

The following information shows the statements and displays used in [Figure E-2](#).

HSC LMUPDEF PARMLIB Member

```
OPTION TITLE('DUAL TCP/IP CONNECTING TO SL3000')
LMUPATH ACS(00) LMUADDR(129.80.71.83,172.27.2.5)
```

TCP/IP Profile Data Set ; OSA CARD #1

```
DEVICE ECCQD01 MPCIPA NONROUTER AUTORESTART
```

```

LINK    ZIPBMVS IPAQENET ECCQD01

; OSA CARD #2
DEVICE ECCQA01 MPCIPA  NONROUTER AUTORESTART
LINK    ZIPB2MVS IPAQENET ECCQA01

HOME
    129.80.16.239 ZIPBMVS
    129.80.65.203 ZIPB2MVS

BEGINROUTES
;
;           NETWORK MASK  FIRSTHOP    LINKNAME    PACKETSIZE
ROUTE 129.80.16.0/24      =          ZIPBMVS     MTU 1492
ROUTE 129.80.65.0/24      =          ZIPB2MVS    MTU 1492
ROUTE 172.27.2.5  HOST    129.80.16.254 ZIPBMVS     MTU 1492
ROUTE 129.80.71.83 HOST    129.80.65.254 ZIPB2MVS    MTU 1492
ROUTE DEFAULT 129.80.16.254          ZIPBMVS     MTU 1492
ROUTE DEFAULT 129.80.65.254          ZIPB2MVS    MTU 1492

ENDROUTES

INCLUDE ZIP.TCPIP.PROFILES(COMMON)

START ECCQD01
START ECCQA01

```

TCP/IP Console Displays

```

D TCPIP,TCPIP,NETSTAT,DEV
EZZ2500I NETSTAT CS V1R4 TCPIP 229
DEVNAME: LOOPBACK          DEVTYPE: LOOPBACK
DEVSTATUS: READY
LNKNAME: LOOPBACK          LNKTYPE: LOOPBACK    LNKSTATUS: READY
NETNUM: 0  QUESIZE: 0
BYTESIN: 1781074          BYTESOUT: 1781074
ACTMTU: 65535
BSD ROUTING PARAMETERS:
MTU SIZE: 00000          METRIC: 00
DESTADDR: 0.0.0.0        SUBNETMASK: 0.0.0.0
MULTICAST SPECIFIC:
MULTICAST CAPABILITY: NO

DEVNAME: ECCQD01          DEVTYPE: MPCIPA
DEVSTATUS: READY          CFGROUTER: NON  ACTROUTER: NON
LNKNAME: ZIPBMVS          LNKTYPE: IPAQENET    LNKSTATUS: READY
NETNUM: 0  QUESIZE: 0    SPEED: 0000001000
BYTESIN: 34704496          BYTESOUT: 11207410
IPBROADCASTCAPABILITY: NO
ARPOFFLOAD: YES  ARPOFFLOADINFO: YES
ACTMTU: 8992
INBPERF: BALANCED
BSD ROUTING PARAMETERS:
MTU SIZE: 00000          METRIC: 00
DESTADDR: 0.0.0.0        SUBNETMASK: 255.255.255.0
MULTICAST SPECIFIC:
MULTICAST CAPABILITY: YES
GROUP                    REFCNT
-----
224.0.0.1                0000000001

DEVNAME: ECCQA01          DEVTYPE: MPCIPA

```

```

DEVSTATUS: READY          CFGROUTER: NON  ACTROUTER: NON
LNKNAME: ZIPB2MVS        LNKTYPE: IPAQENET   LNKSTATUS: READY
NETNUM: 0    QUESIZE: 0   SPEED: 0000000100
BYTESIN: 147508801      BYTESOUT: 2188246
IPBROADCASTCAPABILITY: NO
ARPOFFLOAD: YES  ARPOFFLOADINFO: YES
ACTMTU: 1492
INBPERF: BALANCED
BSD ROUTING PARAMETERS:
MTU SIZE: 00000          METRIC: 00
DESTADDR: 0.0.0.0      SUBNETMASK: 255.255.255.0
MULTICAST SPECIFIC:
MULTICAST CAPABILITY: YES
GROUP          REFCNT
-----
224.0.0.1      0000000001
3 OF 3 RECORDS DISPLAYED
D TCPIP,TCPIP,NETSTAT,ROUTE
EZZ2500I NETSTAT CS V1R4 TCPIP 250
DESTINATION    GATEWAY        FLAGS    REFCNT  INTERFACE
DEFAULT        129.80.16.254  UGS      000003  ZIPBMVS
DEFAULT        129.80.65.254  UGS      000002  ZIPB2MVS
127.0.0.1      0.0.0.0        UH       000003  LOOPBACK
129.80.16.0    0.0.0.0        US       000002  ZIPBMVS
129.80.16.239 0.0.0.0        UH       000000  ZIPBMVS
129.80.65.0    0.0.0.0        US       000000  ZIPB2MVS
129.80.65.203 0.0.0.0        UH       000000  ZIPB2MVS
129.80.71.83   129.80.65.254 UGHS     000001  ZIPB2MVS
172.27.2.5     129.80.16.254 UGHS     000001  ZIPBMVS
9 OF 9 RECORDS DISPLAYED

```

HSC Display Showing SL3000 Connection

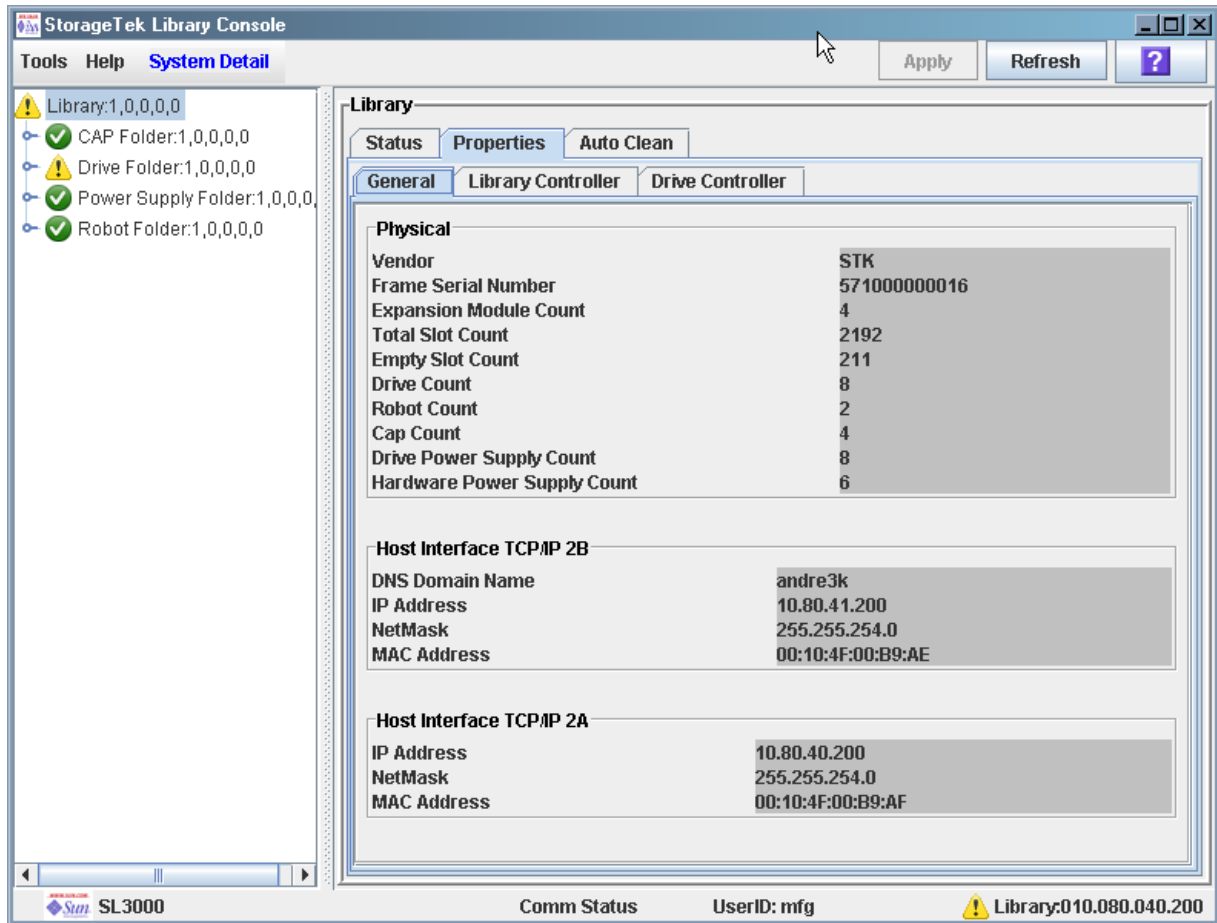
```

D ACS
SLS0000I D ACS
SLS1000I ACS 00 STATUS: CONNECTED 334
ACTIVE QUEUE ELEMENTS 2
COMPATIBILITY LEVELS: HSC=13, LMU=13
DUAL LMU NOT CONFIGURED
IP ADDR 129.80.71.83 ONLINE
IP ADDR 172.27.2.5 ONLINE
SCRATCH VOLUMES AVAILABLE..... 210
FREE CELLS AVAILABLE..... 2008
D TCPIP,TCPIP,NETSTAT,HOME
EZZ2500I NETSTAT CS V1R4 TCPIP 252
HOME ADDRESS LIST:
ADDRESS      LINK          FLG
129.80.16.239 ZIPBMVS      P
129.80.65.203 ZIPB2MVS
127.0.0.1     LOOPBACK
3 OF 3 RECORDS DISPLAYED

```

Figure E-3 shows an SLConsole display of ports 2A and 2B, including IP addresses.

Figure E-3 SLConsole Display of SL3000 Ports 2A/2B IP Addresses

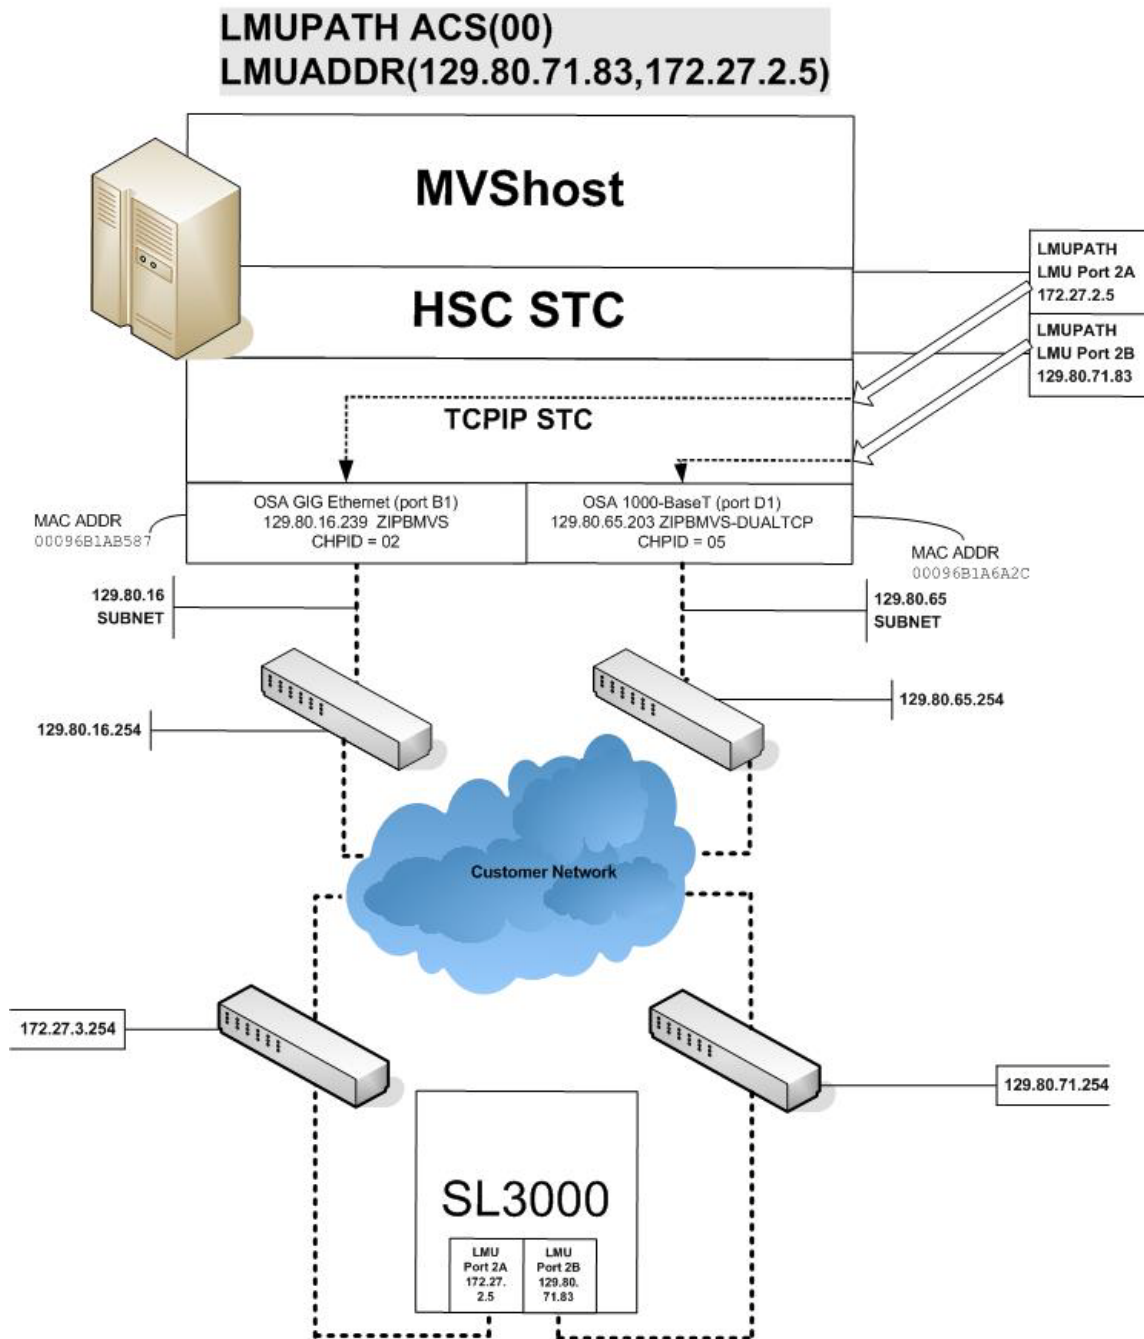


Two SL3000 Network Connections - One Mainframe IP Address

To establish one host IP route from the HSC to two routes to the SL3000, use the process described in this section. To configure the SL3000 dual TCP/IP feature, refer to chapter 4 of the SL3000 Modular Library System Installation Guide.

Figure E-4 shows dual IP (one host IP) with two SL3000 network connections.

Figure E-4 Dual IP (one host IP) with Two SL3000 Network Connections



1. Complete the Network Entries Worksheet (see [Table E-2](#) for a sample) for each set of routes to the SL3000. The worksheet can be found in the *SL3000 Modular Library System Installation Guide*.

Table E-2 Network Entries Worksheet

Port Number	Description	IP Address
2A	HBC Card Port 2A	172.27.2.5
2A	Gateway Port 2A	172.27.3.254

Table E-2 (Cont.) Network Entries Worksheet

Port Number	Description	IP Address
2A	Netmask	/23
2A	Host to Port 2A	129.80.16.239
2B	HBC Card Port 2B	129.80.71.83
2B	Gateway Port 2B	129.80.71.254
2B	Netmask	/23
2B	Host to Port 2B	129.80.16.239

- Optionally, define two dedicated static routes to the SL3000 destination port (2A and 2B) IP addresses over one router.

For example:

```
BEGINROUTES
;   DESTINATION                FIRSTHOP      LINKNAME      PACKETSIZE
ROUTE 129.80.16.0/24           =             &SYSNAME.MVS  MTU 1492
ROUTE 172.27.2.5 HOST         129.80.16.254 &SYSNAME.MVS  MTU 1492
ROUTE 129.80.71.83 HOST      129.80.16.254 &SYSNAME.MVS  MTU 1492
ROUTE DEFAULT                 129.80.16.254 &SYSNAME.MVS  MTU 1492

ENDROUTES
```

- Define a second LMUADDR parameter for the port 2A IP address on the SL3000.

For example:

```
LMUPATH ACS(00) LMUADDR(129.80.71.83,172.27.2.5)
```

- Enter the LMUPDEF command containing the LMUPATH statements that define the host name or IP address for each ACS.

```
LMUPDEF DSN('xxx.xxx.xxx(xxx)')
```

In the following example, LMUPDEF loads LMUPATH parameters from YOUR.DSN(MEMBER).

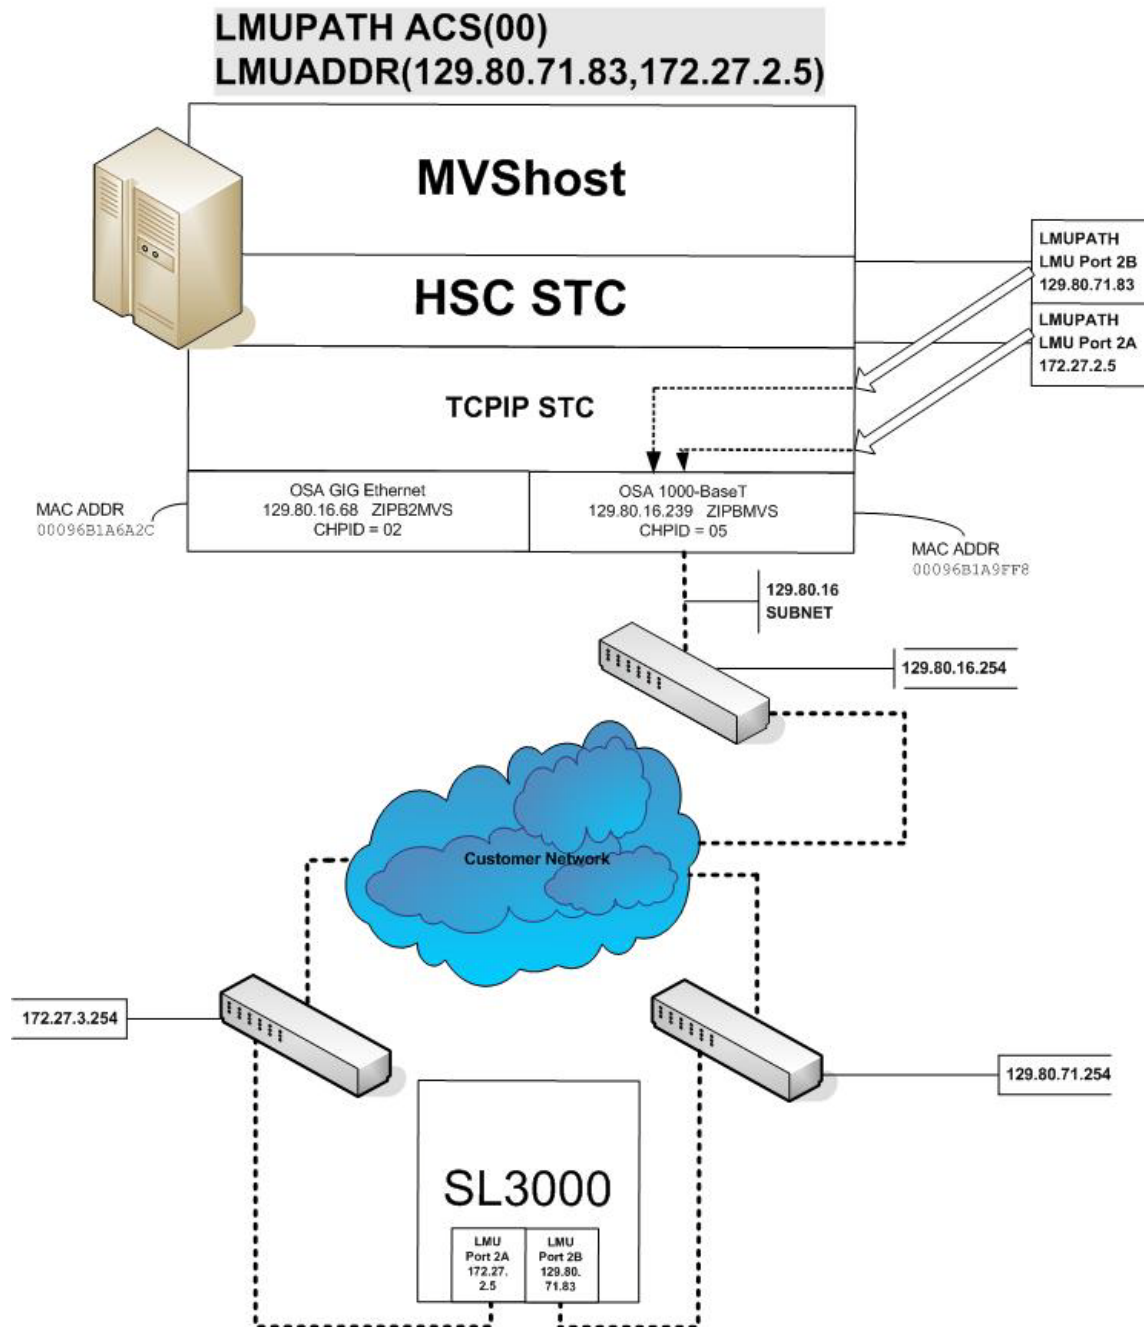
```
LMUPDEF DSN('YOUR.DSN(MEMBER)')
```

- Allow the trained SL3000 service representative to enter the network connections to the SL3000 library for either port 2A and 2B, whichever is applicable.
- Vary the ACS offline and back online to pick up the revised LMUPATH statement that includes the second connection. This can be done one host at a time to minimize down time.

Sample Configuration - One Host IP, Two SL3000 Network Connections

Figure E-5 shows a sample dual IP (one host IP) configuration with two SL3000 network connections, followed by the statements that apply to it.

Figure E-5 Dual IP (one host IP) with Two SL3000 Network Connections



Setup Statements and Displays

The following information shows the statements and displays used in [Figure E-5](#).

HSC LMUPDEF PARMLIB Member

```
OPTION TITLE('DUAL TCPIP CONNECTING TO SL30001')
LMUPATH ACS(00) LMUADDR(129.80.71.83,172.27.2.5)
```

TCP/IP Profile Data Set

```
; OSA CARD #1
DEVICE ECCQD01 MPCIPA NONROUTER AUTORESTART
```

```

LINK    ZIPBMVS IPAQENET ECCQD01

HOME
  129.80.&IPADDR1 &SYSNAME.MVS

BEGINROUTES
;
      NETWORK MASK  FIRSTHOP      LINKNAME      PACKETSIZE
ROUTE 129.80.16.0/24      =          &SYSNAME.MVS  MTU 1492
ROUTE 172.27.2.5 HOST    129.80.16.254 &SYSNAME.MVS  MTU 1492
ROUTE 129.80.71.83 HOST  129.80.16.254 &SYSNAME.MVS  MTU 1492
ROUTE DEFAULT

ENDROUTES

INCLUDE ZIP.TCPIP.PROFILES(COMMON)

START ECCQD01

```

TCP/IP Console Displays

```

D TCPIP,,N,DEV
EZZ2500I NETSTAT CS V1R4 TCPIP 931
DEVNAME: LOOPBACK          DEVTYPE: LOOPBACK
DEVSTATUS: READY
LNKNAME: LOOPBACK          LNKTYPE: LOOPBACK    LNKSTATUS: READY
NETNUM: 0  QUESIZE: 0
BYTESIN: 2136824          BYTESOUT: 2136824
ACTMTU: 65535
BSD ROUTING PARAMETERS:
  MTU SIZE: 00000          METRIC: 00
  DESTADDR: 0.0.0.0        SUBNETMASK: 0.0.0.0
MULTICAST SPECIFIC:
  MULTICAST CAPABILITY: NO
DEVNAME: ECCQD01          DEVTYPE: MPCIPA
DEVSTATUS: READY          CFGROUTER: NON  ACTROUTER: NON
LNKNAME: ZIPBMVS          LNKTYPE: IPAQENET  LNKSTATUS: READY
BYTESIN: 48605838          BYTESOUT: 9790950
NETNUM: 0  QUESIZE: 0  SPEED: 0000001000
IPBROADCASTCAPABILITY: NO
ARPOFFLOAD: YES  ARPOFFLOADINFO: YES
ACTMTU: 8992
INBPERF: BALANCED
BSD ROUTING PARAMETERS:
  MTU SIZE: 00000          METRIC: 00
  DESTADDR: 0.0.0.0        SUBNETMASK: 255.255.255.0
MULTICAST SPECIFIC:
  MULTICAST CAPABILITY: YES
  GROUP          REFCNT
  -----          -----
  224.0.0.1          0000000001
2 OF 2 RECORDS DISPLAYED

```

```

D TCPIP,,N,ROUTE
EZZ2500I NETSTAT CS V1R4 TCPIP 933
DESTINATION      GATEWAY          FLAGS      REFCNT  INTERFACE
DEFAULT          129.80.16.254    UGS        000001  ZIPBMVS
127.0.0.1        0.0.0.0          UH         000004  LOOPBACK
129.80.16.0      0.0.0.0          US         000000  ZIPBMVS
129.80.16.239    0.0.0.0          UH         000000  ZIPBMVS
129.80.71.83     129.80.16.254    UGHS       000001  ZIPBMVS
172.27.2.5       129.80.16.254    UGHS       000001  ZIPBMVS

```

```

6 OF 6 RECORDS DISPLAYED

D TCPIP,,N,HOME
EZZ2500I NETSTAT CS V1R4 TCPIP 935
HOME ADDRESS LIST:
ADDRESS          LINK          FLG
129.80.16.239   ZIPBMVS      P
127.0.0.1       LOOPBACK
2 OF 2 RECORDS DISPLAYED
    
```

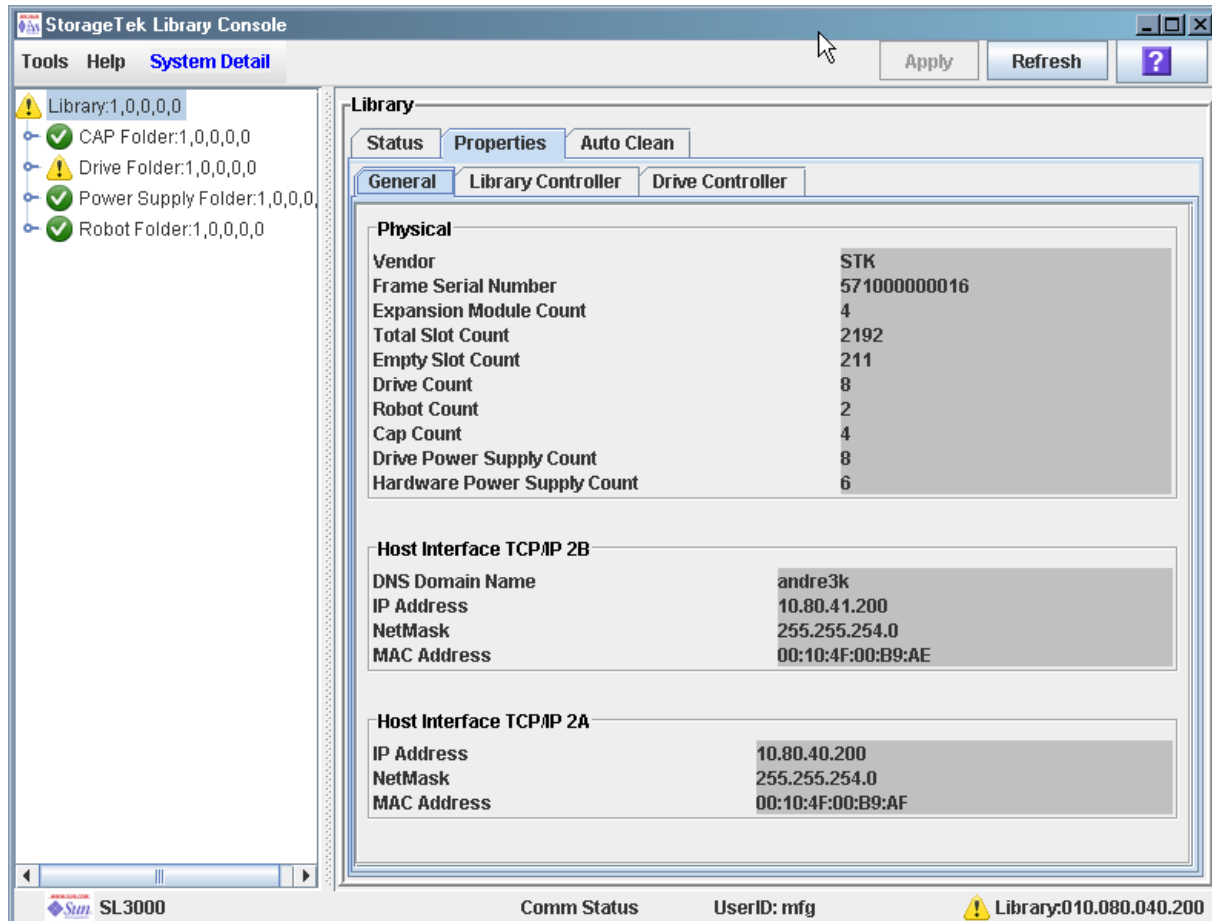
HSC Display Showing SL3000 Connection

```

D ACS
SLS0000I D ACS
SLS1000I ACS 00 STATUS: CONNECTED 942
ACTIVE QUEUE ELEMENTS          1
COMPATIBILITY LEVELS: HSC=13, LMU=13
DUAL LMU NOT CONFIGURED
IP ADDR 129.80.71.83 ONLINE
IP ADDR 172.27.2.5 ONLINE
SCRATCH VOLUMES AVAILABLE..... 210
FREE CELLS AVAILABLE.....      2007
    
```

Figure E-6 shows an SLConsole display of ports 2A and 2B, including IP addresses.

Figure E-6 SLConsole Display of SL3000 Ports 2A/2B IP Addresses



Troubleshooting Library Connectivity Problems

For more information, see "[Troubleshooting Library Connectivity Problems](#)."

VSM4 ESCON Configuration

VSM4 provides the following advantages over its predecessors:

- Enhanced connectivity options
- Greater throughput
- Greater VTSS capacity
- 4x the number of VTDs and 3x the maximum number of VTVs per VTSS
- Improved reliability and serviceability

[Table F-1](#) summarizes the VSM3 to VSM4 ESCON enhancements that you see from a software and system configuration perspective.

Table F-1 VSM3 to VSM4 Comparison: Software and System Configuration ESCON Enhancements

Product Feature	VSM3	VSM4
ESCON Interfaces	16 total where: <ul style="list-style-type: none"> ■ 2 to 14 can be host channels ■ 2 to 8 can be Nearlink/CLINK connections 	32 total where: <ul style="list-style-type: none"> ■ 2 to 28 can be host channels ■ 2 to 16 can be Nearlink/CLINK connections <p>Note: VSM4s are shipped with 16 ports enabled. With the 16 ports enabled option, only the top port on each CIP is enabled (Port 0 or Port 2). 32 ports enabled is an optional, separately priced feature that is activated through the microcode diskette.</p> <p>On a VSM4 with 32 ports enabled, each ICE3 ESCON interface card contains two pairs of ESCON ports.</p>

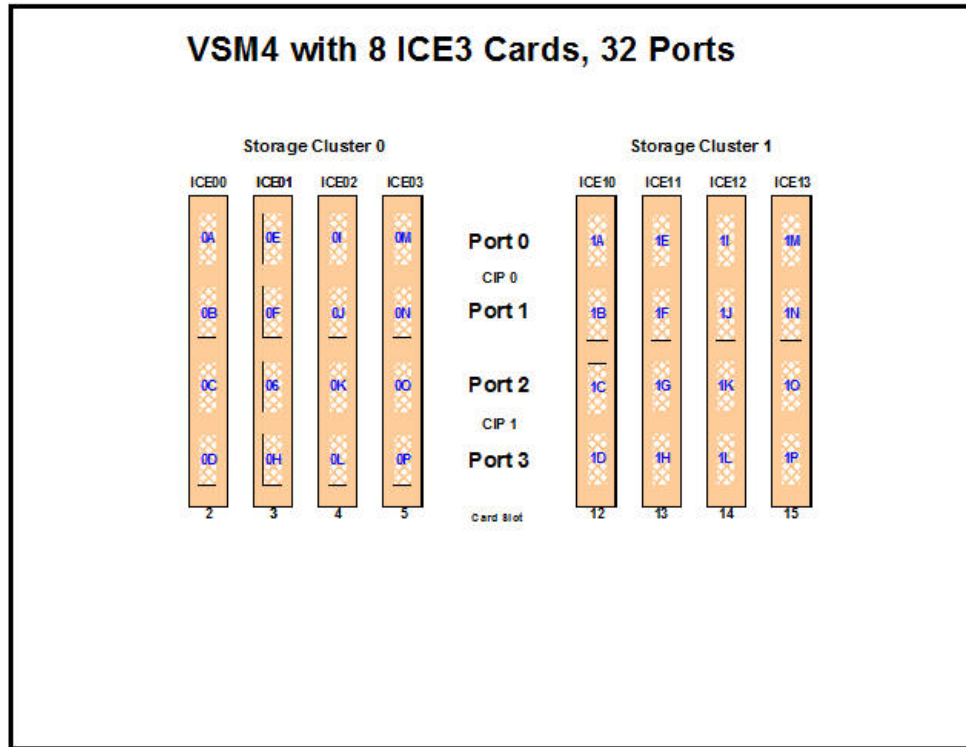
Table F-1 (Cont.) VSM3 to VSM4 Comparison: Software and System Configuration ESCON Enhancements

Product Feature	VSM3	VSM4
Maximum Logical Paths	128	16 per port for the 16 port standard configuration = 256 logical paths 16 per port for the 32 port optional configuration= 512 logical paths Note: VSM4 provides a theoretical maximum of 512 logical paths per VTSS but you cannot allocate all 512 logical paths for host-to-VTSS connections.
VTDs per VTSS	64	256
Maximum resident VTVs per VTSS	100,000	300,000

VSM4 with 32 Ports

For the 32 port option, the 8 ICE3 cards have four ESCON ports per card as shown in [Figure F-1](#).

Figure F-1 VSM4 with 32 Ports



Note: In [Figure F-1](#) and all the other figures in this appendix, the ports are shown with their channel interface identifiers for **enabled** ports (32 in [Figure F-1](#)). These channel interface identifiers are the values that are required for the CHANIF values that you code for the CONFIG utility. Each value is two characters in length and has a value from 0A to 1P. The first digit is the VTSS cluster ID (valid values are 0 or 1). The second digit is the group or adapter ID (valid values are A to P).

In [Figure F-1](#), note the following:

- Each ICE3 card contains two pairs of ESCON ports. Each pair is controlled by its own Channel Interface Processor (CIP).
- With microcode level D02.06.00.00 or higher, multiple Nearlink device connections using an ESCON switch or Director on the same port now allow:
 - **Up to a total of 16 simultaneous NearLink I/O transfers**, which can be spread across multiple targets on as many as 14 NearLink ports.
 - **Up to a total of 2 simultaneous NearLink I/O transfers are allowed per port.**
- For a VSM4, each CIP can operate with only one of two modes, which is set at the VTSS LOP:
 - **Host Mode.** In Host Mode, either or both ESCON ports can connect to host CPU channels, including through ESCON Director(s) or channel extenders. Ports of a CIP in Host mode cannot connect to RTDs or to Secondary VTSSs using CLINKS. Note, however, that Secondary VTSSs must have an ESCON port in Host Mode to connect using a CLINK from an ESCON port in Nearlink Mode in a Primary VTSS.

Also note that you can have two physical paths from the same LPAR to the same CIP, as long as the two physical paths address different (not overlapping) logical control units. For example, a single host LPAR can address logical control units 0-7 on one CIP port, and 8-F on the other CIP port of the same CIP.
 - **Nearlink Mode.** In Nearlink Mode, either or both ESCON ports can connect to an RTD or using a CLINK to a Secondary VTSS. Ports of a CIP in Nearlink mode cannot connect to host CPU channels. You can set a maximum of 8 CIPs to Nearlink Mode, and here's the important discussion: only one Nearlink port per CIP is active at one time.

[Table F-2](#) shows best practices for optimizing port operations.

Table F-2 Optimizing VSM4 Port Operations

Configuration - Two Ports on a CIP	Best Practices
Two CLINKs	Attach a maximum of 2 because each port allows two active operations. Note, however, that these operations share the bandwidth of the port.
CLINK and RTD	An advantage if you attach one CLINK originator/one RTD per director, because both can be active.

Table F–2 (Cont.) Optimizing VSM4 Port Operations

Configuration - Two Ports on a CIP	Best Practices
Two RTDs	<p>An advantage for the following:</p> <ul style="list-style-type: none"> ■ Optimize use of local and remote RTDs. During busy shifts, use only local RTDs on the FIP. During quiet periods, switch to remote RTDs for deep archive and DR work. Because you can have two active devices, you can also simultaneously run one local and one remote RTD. Note, however, that these operations share the bandwidth of the port. ■ Optimize use of different drive technologies. As described in the previous bullet, use a T9840 as a local RTD, then switch to a T9940 for deep archive. You can also use this feature to migrate from older drive technology (such as 9490) to newer technology (such as 9840). Use Management and Storage Classes to read in data from older media, then switch to the newer technology drive to place data on new media. This technique effectively gives you greater physical connectivity to different drive technologies without incurring the overhead of full time, real time FICON connections to each drive type. Also as above, because you can have two active devices, you can also simultaneously run two RTDs with different drive technologies. Note, however, that these operations share the bandwidth of the port.

- On a VSM4 with 32 ports enabled, you have a theoretical maximum of 512 logical paths on the VSM4. However, you must have some RTD connections so you cannot allocate all 512 logical paths for host-to-VTSS connections. The minimum number of RTDs is determined by a couple factors:
 - CONFIG will not allow fewer than 2 RTDS per VTSS.
 - CONFIG cannot check the device type but Oracle strongly recommends at least two RTDS of each device type in each ACS to which the VTSS is attached. Otherwise, you can seriously compromise error recovery, and also impact the efficiency of space reclamation.

If you have only two RTDs, Best Practices suggest that you connect them to different ICE3 cards. Once you have done that, you have effectively used up 4 Nearlink ports due to the CIP mode nature of the ICE3 card. Therefore, in an 8 ICE card configuration, this leaves 28 available ports for host-to-VTSS ESCON channel connections, which equals a maximum of 16 x 28 (448) logical paths. For more information, see "[Logical Paths for VSM 4 with 32 Ports.](#)"
- A host logical path is the communication path between a host and all of the 256 VTDs within the VSM4. [Table F–3](#) summarizes the configuration options and maximum host logical paths for a VSM4 with 32 enabled ports.

Table F–3 VSM4 Configuration Options - 32 Ports

Host CIPs	Maximum Host Connections	Nearlink CIPs	Max Nearlink Connections	Maximum Host Logical Paths
8	16	8	16	256
9	18	7	14	288
10	20	6	12	320
11	22	5	10	352
12	24	4	8	384

Table F-3 (Cont.) VSM4 Configuration Options - 32 Ports

Host CIPs	Maximum Host Connections	Nearlink CIPs	Max Nearlink Connections	Maximum Host Logical Paths
14	28	2	4	448

- In HCD:
 - From a single MVS host, you can only define 8 channels (CHPIDs) running to a single control unit (single VSM4). Also note that ICE3 cards **cannot** have 2 paths from the same LPAR connected to two ports with a common CIP.
 - You use the CNTLUNIT statement to define each VSM4 as 16 3490 images.
 - You use the IODEVICE statement to define the 16 VTDs that are associated with each 3490 image.

VSM4 Configuration Examples - 32 Ports

For VSM4s with 32 ports, look at two examples of port configurations:


- ["VSM4 Configuration Example: 16 Host Ports, 16 RTD Ports"](#)
- ["VSM 4 Configuration Example: 20 Host Ports, 12 RTD Ports"](#)

For a VSM4 host gen example, see ["IOCP Example for Single MVS Host Connected to a VSM4 Using ESCON Directors."](#)

VSM4 Configuration Example: 16 Host Ports, 16 RTD Ports

[Figure F-2](#) shows CONFIG channel interface identifiers of 16 for hosts, 16 for RTDs for a VSM4.

Figure F-2 VSM4 with 16 Host Ports, 16 RTD Ports




CONFIG Example for VSM4 with 16 Host Ports, 16 RTD Ports

```
//CREATECF EXEC PGM=SLUADMIN, PARM='MIXED'
//STEPLIB DD DSN=hlq.SEALINK, DISP=SHR
//SLSCNTL DD DSN=hlq.DBASEPRM, DISP=SHR
//SLSCNTL2 DD DSN=hlq.DBASESEC, DISP=SHR
//SLSSTBY DD DSN=hlq.DBASETBY, DISP=SHR
//SLSPRINT DD SYSOUT=*
//SLSIN DD *
CONFIG
GLOBAL MAXVTV=32000 MVCFREE=40
RECLAIM THRESHLD=70 MAXMVC=40 START=35
VTSS NAME=VSM401 LOW=70 HIGH=80 MAXMIG=8 RETAIN=5
RTD NAME=VSM42A02 DEVNO=2A02 CHANIF=0G
RTD NAME=VSM42A04 DEVNO=2A04 CHANIF=0K
RTD NAME=VSM42A06 DEVNO=2A06 CHANIF=0O
RTD NAME=VSM42A0A DEVNO=2A0A CHANIF=1G
RTD NAME=VSM42A0C DEVNO=2A0C CHANIF=1K
RTD NAME=VSM42A0E DEVNO=2A0E CHANIF=1O
VTD LOW=9900 HIGH=99FF
```

VSM 4 Configuration Example: 20 Host Ports, 12 RTD Ports

Figure F-3 shows port assignments of 20 for hosts, 12 for RTDs for a VSM4.

Figure F-3 VSM4 with 20 Host Ports, 12 RTD Ports

**CONFIG Example for VSM4 with 20 Host Ports, 12 RTD Ports**


```

//CREATECF EXEC PGM=SLUADMIN, PARM='MIXED'
//STEPLIB DD DSN=hlq.SEALINK, DISP=SHR
//SLSCNTL DD DSN=hlq.DBASEPRM, DISP=SHR
//SLSCNTL2 DD DSN=hlq.DBASESEC, DISP=SHR
//SLSSTBY DD DSN=hlq.DBASETBY, DISP=SHR
//SLSPRINT DD SYSOUT=*
//SLSIN DD *
CONFIG
  GLOBAL MAXVTV=32000 MVCFREE=40
  RECLAIM THRESHLD=70 MAXMVC=40 START=35
  VTSS NAME=VSM401 LOW=70 HIGH=80 MAXMIG=6 RETAIN=5
  RTD NAME=VSM42A00 DEVNO=2A00 CHANIF=0G
  RTD NAME=VSM42A01 DEVNO=2A01 CHANIF=0H
  RTD NAME=VSM42A02 DEVNO=2A02 CHANIF=0K
  RTD NAME=VSM42A03 DEVNO=2A03 CHANIF=0L
  RTD NAME=VSM42A04 DEVNO=2A04 CHANIF=0O
  RTD NAME=VSM42A05 DEVNO=2A05 CHANIF=0P
  RTD NAME=VSM42A06 DEVNO=2A06 CHANIF=1G
  RTD NAME=VSM42A07 DEVNO=2A07 CHANIF=1H
  RTD NAME=VSM42A08 DEVNO=2A08 CHANIF=1K
  RTD NAME=VSM42A09 DEVNO=2A09 CHANIF=1L
  RTD NAME=VSM42A0A DEVNO=2A0A CHANIF=1O
  RTD NAME=VSM42A0B DEVNO=2A0B CHANIF=1P
  VTD LOW=9900 HIGH=99FF
  
```

VSM4 with 16 Ports

For the 16 port option, the 8 ICE3 cards have two ESCON ports per card as shown in [Figure F-4](#).

Figure F-4 VSM4 with 16 Ports



Note: In Figure F-4 and all the other figures in this appendix, the ports are shown with their channel interface identifiers for enabled ports (16 in Figure F-4). These channel interface identifiers are the values that are required for the CHANIF values that you code for the CONFIG utility. Each value is two characters in length and has a value from 0A to 1P. The first digit is the VTSS cluster ID (valid values are 0 or 1). The second digit is the group or adapter ID (valid values are A to P).

In Figure F-4, note the following:

- Each ICE3 card has two CIPs with a single port enabled on each CIP. As with the 32 port option, each CIP can operate with only one of two modes, which is set at the VTSS LOP:
 - **Host Mode.** In Host Mode, the single ESCON port can connect to host CPU channels, including through ESCON Director(s) or channel extenders. Ports of a CIP in Host mode cannot connect to RTDs or to Secondary VTSSs using CLINKS. Note, however, that Secondary VTSSs must have an ESCON port in Host Mode to connect through a CLINK from an ESCON port in Nearlink Mode in a Primary VTSS.
 - **Nearlink Mode.** In Nearlink Mode, the single ESCON port can connect to an RTD or through a CLINK to a Secondary VTSS. Ports of a CIP in Nearlink mode cannot connect to host CPU channels.

You can set a maximum of 8 CIPs to Nearlink Mode. Therefore, in a 16 port configuration, the single port on a CIP can be either a CLINK or an RTD connection.

- On a VSM4 with 16 ports enabled, you have a theoretical maximum of 256 logical paths on the VSM4. However, you must have some RTD connections so you cannot allocate all 256 logical paths for host-to-VTSS connections.

What is the minimum number of RTDs?

- CONFIG will not allow fewer than 2 RTDS per VTSS.
- CONFIG cannot check device type but Oracle strongly recommends at least two RTDS of each device type in each ACS to which the VTSS is attached. Otherwise, you can seriously compromise error recovery and also impact the efficiency of space reclamation. If you had only two RTDs, Best Practices would suggest that you connect them to different ICE3 cards. Once you have done that, you have effectively used up 4 Nearlink ports. Therefore, in an 8 ICE card 16 port configuration, this leaves 12 available ports for host-to-VTSS ESCON channel connections, which equals a maximum of 16 x 12 or 192 logical paths. For more information, see "[Logical Paths for VSM 4 with 32 Ports.](#)"
- A host logical path is the communication path between a host and all of the 256 VTDs within the VSM4. [Table F-4](#) summarizes the configuration options and maximum host logical paths for a VSM4 with 16 enabled ports.

Table F-4 VSM4 Configuration Options - 16 Ports

Host CIPs	Maximum Host Connections	Nearlink CIPs	Max Nearlink Connections	Maximum Host Logical Paths
8	16	8	8	128
9	18	7	7	144
10	20	6	6	160
11	22	5	5	176
12	24	4	4	192
14	28	2	2	224

- In HCD:
 - From a single MVS host, you can only define 8 channels (CHPIDs) running to a single control unit (single VSM4). Also note that ICE3 cards cannot have 2 paths from the same LPAR connected to two ports with a common CIP.
 - You use the CNTLUNIT statement to define each VSM4 as 16 3490 images.
 - You use the IODEVICE statement to define the 16 VTDs that are associated with each 3490 image.

VSM4 Configuration Examples - 16 Ports

For VSM4s with 16 ports, look at two examples of port configurations:

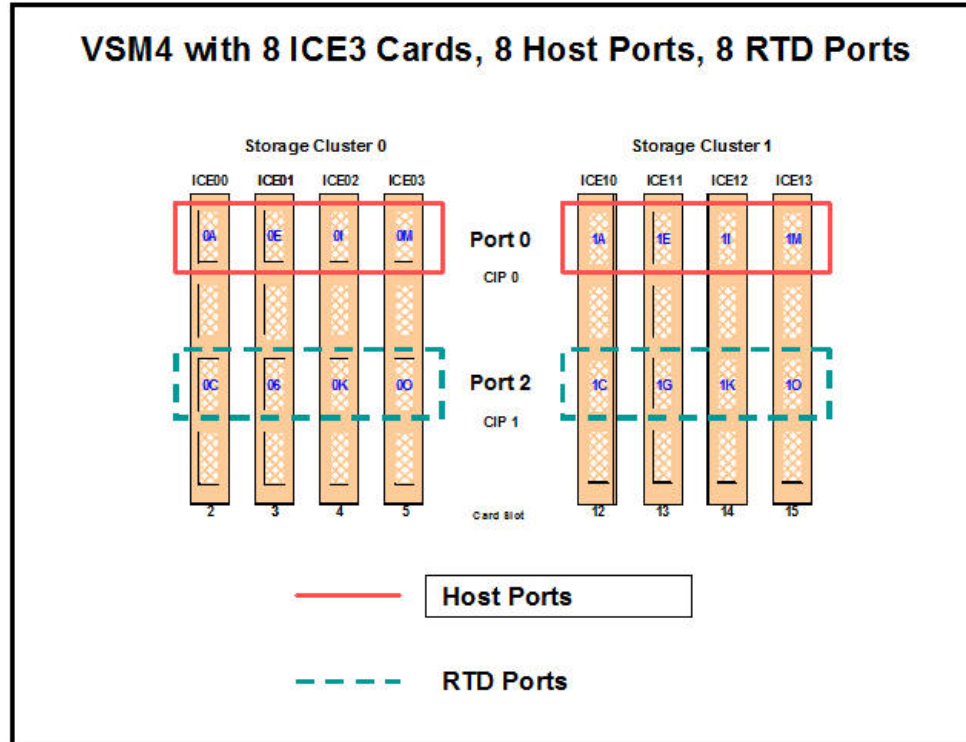
- "[VSM4 Configuration Example: 8 Host Ports, 8 RTD Ports](#)"
- "[VSM 4 Configuration Example: 10 Host Ports, 6 RTD Ports](#)"

For a VSM4 host gen example, see "[IOCP Example for Single MVS Host Connected to a VSM4 Using ESCON Directors.](#)"

VSM4 Configuration Example: 8 Host Ports, 8 RTD Ports

Figure F-5 shows CONFIG channel interface identifiers of 8 for hosts, 8 for RTDs for an 8 ICE3 card VSM4 with 16 ports.

Figure F-5 VSM4 with 8 Host Ports, 8 RTD Ports



CONFIG Example for VSM4 with 8 Host Ports, 8 RTD Ports

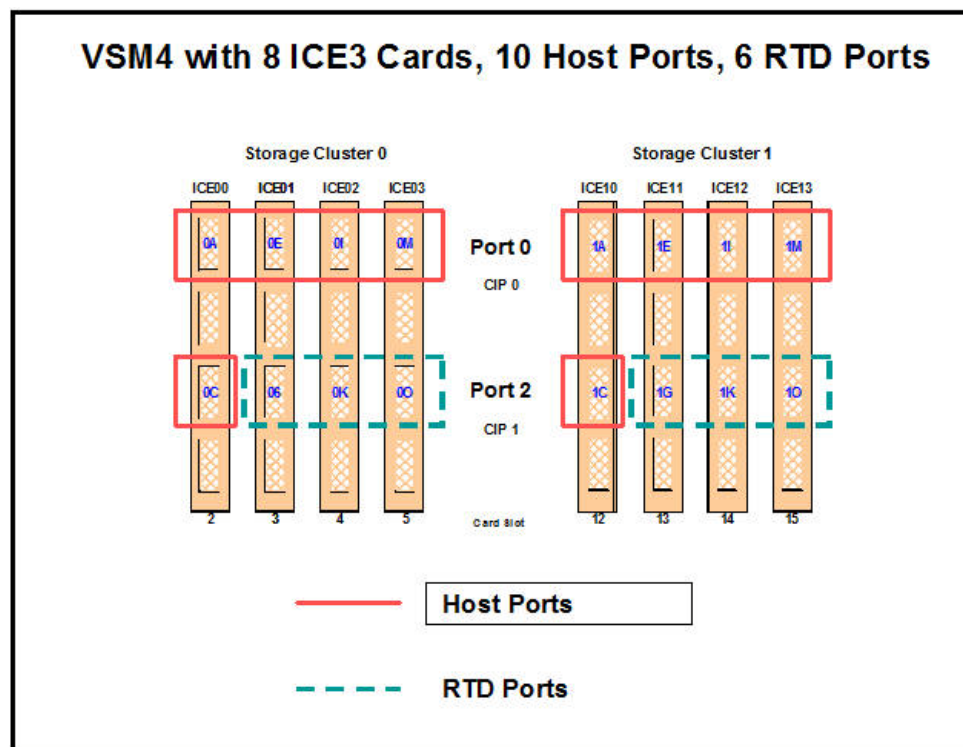
The following example shows CONFIG JCL to define the VSM4 configuration shown in Figure F-5.

```
//CREATECF EXEC PGM=SLUADMIN, PARM='MIXED'
//STEPLIB DD DSN=hlq.SEALINK, DISP=SHR
//SLSCNTL DD DSN=hlq.DBASEPRM, DISP=SHR
//SLSCNTL2 DD DSN=hlq.DBASESEC, DISP=SHR
//SLSSTBY DD DSN=hlq.DBASETBY, DISP=SHR
//SLSPRINT DD SYSOUT=*
//SLSIN DD *
CONFIG
GLOBAL MAXVTV=32000 MCVFREE=40
RECLAIM THRESHLD=70 MAXMVC=40 START=35
VTSS NAME=VSM401 LOW=70 HIGH=80 MAXMIG=8 RETAIN=5
RTD NAME=VSM42A00 DEVNO=2A00 CHANIF=0C
RTD NAME=VSM42A02 DEVNO=2A02 CHANIF=0G
RTD NAME=VSM42A04 DEVNO=2A04 CHANIF=0K
RTD NAME=VSM42A06 DEVNO=2A06 CHANIF=0O
RTD NAME=VSM42A08 DEVNO=2A08 CHANIF=1C
RTD NAME=VSM42A0A DEVNO=2A0A CHANIF=1G
RTD NAME=VSM42A0C DEVNO=2A0C CHANIF=1K
RTD NAME=VSM42A0E DEVNO=2A0E CHANIF=1O VTD LOW=9900 HIGH=99FF
```

VSM 4 Configuration Example: 10 Host Ports, 6 RTD Ports

Figure F-6 shows port assignments of 10 for hosts, 6 for RTDs for a VSM4 with 16 ports.

Figure F-6 VSM4 with 10 Host Ports, 6 RTD Ports



CONFIG Example for VSM4 with 10 Host Ports, 6 RTD Ports

The following example shows CONFIG JCL to define the VSM4 configuration shown in Figure F-6.


```
//CREATECF EXEC PGM=SLUADMIN, PARM='MIXED'
//STEPLIB DD DSN=hlq.SEALINK, DISP=SHR
//SLSCNTL DD DSN=hlq.DBASEPRM, DISP=SHR
//SLSCNTL2 DD DSN=hlq.DBASESEC, DISP=SHR
//SLSSTBY DD DSN=hlq.DBASESTBY, DISP=SHR
//SLSPRINT DD SYSOUT=*
//SLSIN DD *
CONFIG
GLOBAL MAXVTV=32000 MVCFREE=40
RECLAIM THRESHLD=70 MAXMVC=40 START=35
VTSS NAME=VSM401 LOW=70 HIGH=80 MAXMIG=8 RETAIN=5
RTD NAME=VSM42A02 DEVNO=2A02 CHANIF=0G
RTD NAME=VSM42A04 DEVNO=2A04 CHANIF=0K
RTD NAME=VSM42A06 DEVNO=2A06 CHANIF=0O
RTD NAME=VSM42A0A DEVNO=2A0A CHANIF=1G
RTD NAME=VSM42A0C DEVNO=2A0C CHANIF=1K
RTD NAME=VSM42A0E DEVNO=2A0E CHANIF=1O
VTD LOW=9900 HIGH=99FF
```

IOCP Example for Single MVS Host Connected to a VSM4 Using ESCON Directors

Figure F-7 shows a configuration diagram for a single MVS host connected to a VSM4 using ESCON Directors, and the example following it contains IOCP statements for this configuration. Note that:

- From MVSA, you define 8 CHPIDs, with each path switched in the ESCON Director, for a total of 8 channels running to the VSM4.
- You code 16 CNTLUNIT statements to define the VSM4 as 16 3490 images.
- You code IODEVICE statement to define the 16 VTDs that are associated with each 3490 image.

Figure F-7 Configuration Diagram: Single MVS Host Connected to a VSM4 Using ESCON Directors



```
ESCD4C CHPID PATH=(20,70),TYPE=CNC,SWITCH=4C
ESCD4D CHPID PATH=(21,71),TYPE=CNC,SWITCH=4D
ESCD4E CHPID PATH=(30,80),TYPE=CNC,SWITCH=4E
ESCD4F CHPID PATH=(31,81),TYPE=CNC,SWITCH=4F
```

```
CU1 CNTLUNIT CUNUMBR=001,
    PATH=(20,21,30,31,70,71,80,81),
    LINK=(D0,D4,D0,D4,D4,D0,D4,D0),
    UNIT=3490,CUADD=0,
    UNITADD=((00,16))
```

```
STRING1 IODEVICE ADDRESS=(0500,16),
    CUNUMBR=(001),
    UNIT=3490,
    UNITADD=00,STADET=Y
```

```
CU2 CNTLUNIT CUNUMBR=002,
    PATH=(20,21,30,31,70,71,80,81),
    LINK=(D0,D4,D0,D4,D4,D0,D4,D0),
    UNIT=3490,CUADD=1,
```

```

UNITADD= ( (00,16) )

STRING2  IODEVICE ADDRESS=(0510,16),
         CUNUMBER=(002),
         UNIT=3490,
         UNITADD=00,STADET=Y
         .
         .
         .

CU15 CNLUNIT CUNUMBR=015,
     PATH=(20,21,30,31,70,71,80,81),
     LINK=(D0,D4,D0,D4,D4,D0,D4,D0),
     UNIT=3490,CUADD=E,
     UNITADD=( (00,16) )

STRING15 IODEVICE ADDRESS=(05E0,16),
        CUNUMBER=(015),
        UNIT=3490,
        UNITADD=00,STADET=Y

CU16 CNLUNIT CUNUMBR=016,
     PATH=(20,21,30,31,70,71,80,81),
     LINK=(D0,D4,D0,D4,D4,D0,D4,D0),
     UNIT=3490,CUADD=F,
     UNITADD=( (00,16) )

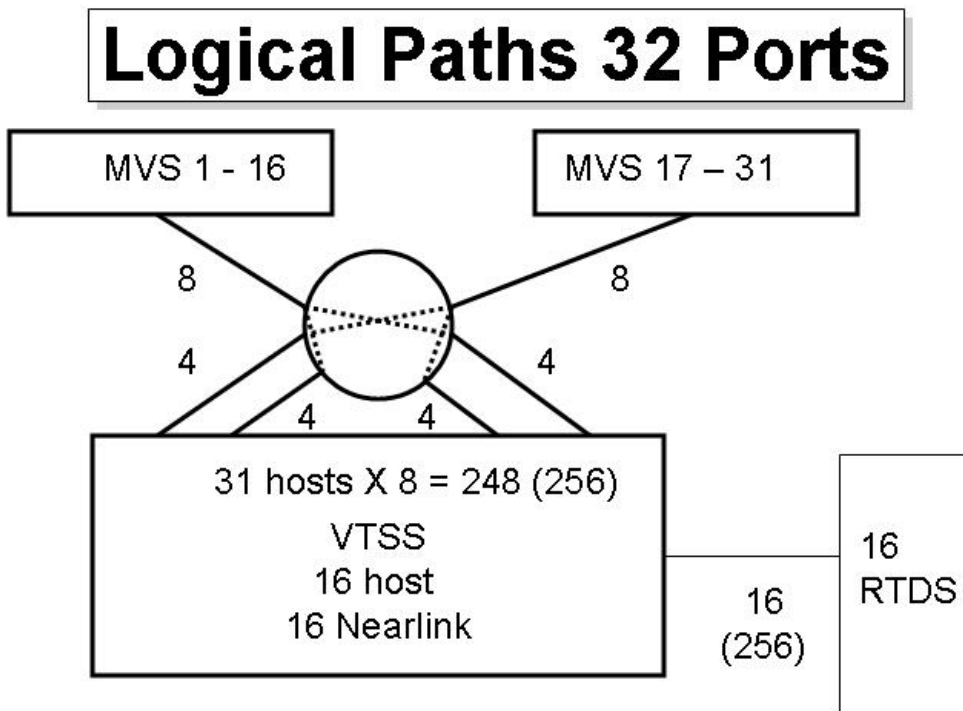
STRING16 IODEVICE ADDRESS=(05F0,16),
        CUNUMBER=(016),
        UNIT=3490,
        UNITADD=00,STADET=Y

```

Logical Paths for VSM 4 with 32 Ports

A VSM4 with 32 ports has 4x the number of logical paths available to VSM2s and VSM3s. Does this mean that a VSM4 has enough logical paths for connectivity, redundancy, and throughput for all attached hosts? Even with 16 RTDs and 31 hosts attached, the answer is "yes" as shown in [Figure F-8](#).

Figure F-8 Logical Paths for VSM 4 with 32 Ports, 31 Hosts, 16 RTDs



In [Figure F-8](#):

- The 16 RTDs consume 16 x 16 or 256 logical paths.
- The **maximum** logical paths you allocated for a VSM2/3 was 4 to a host requiring maximum throughput (which also satisfied the redundancy/connectivity requirements). Therefore, if you allocated double that number, or 8 logical paths, for each of the 31 hosts in this configuration, you only consume 248, or 8 less than the logical paths remaining for host connections.

Therefore, logical path allocation is not an issue, as it was with VSM2s and VSM3s.

VSM4 FICON Front-End and Back-End Configuration

The VSM4 FICON Back-End connectivity feature adds value to the previously available FICON front-end connectivity. [Table G-1](#) summarizes the supported card configurations for VSM4 FICON Front-End plus Back-End connectivity.

Table G-1 Supported Card Configurations for VSM4 FICON Front-End plus Back-End Connectivity

VCF Cards	FICON Ports	ICE Cards	ESCON Ports	Total Ports	Total Logical Paths (16 per ICE Port, 64 per VCF Port)
2	4	6	24	28	640
4	8	4	16	24	768
6	12	2	8	20	896
8	16	0	0	16	1024

VSM4 FICON VCF Card Options

VSM4 supports the following FICON VCF card options:

- [Figure G-1](#) shows a VSM4 with 6 ICE cards, 2 VCF cards.
- [Figure G-2](#) shows a VSM4 with 4 ICE cards, 4 VCF cards.
- [Figure G-3](#) shows a VSM4 with 2 ICE cards, 6 VCF cards.
- [Figure G-4](#) shows a VSM4 with 8 VCF cards.

Figure G-1 VSM4 with 6 ICE Cards, 2 VCF Cards




Figure G-2 VSM4 with 4 ICE Cards, 4 VCF Cards




Figure G-3 VSM4 with 2 ICE Cards, 6 VCF Cards





Figure G-4 VSM4 with 8 VCF Cards



Note: In [Figure G-1](#) through [Figure G-4](#), the VCF cards must go in:

- Slots 5 and 15 in a two-VCF card configuration
 - Slots 3, 5, 13, and 15 in a four-VCF card configuration.
 - Slots 3, 4, 5, 13, 14, and 15 in a six-VCF card configuration.
 - All slots in an eight-VCF card configuration.
-
-

- FICON ports are controlled by a FICON Interface processor (FIP), ESCON ports are controlled by a CIP. Regardless of the card configuration, there can be only a total of 14 Nearlink FIPs and/or CIPs.

With microcode level D02.06.00.00 or higher, multiple Nearlink device connections through a FICON or ESCON switch or Director on the same port now allow:

- **Up to a total of 16 simultaneous NearLink I/O transfers**, which can be spread across multiple targets on as many as 14 NearLink ports.
- **Up to a total of 2 simultaneous NearLink I/O transfers are allowed per port.**
- All FICON ports can be configured as either a Host port or Nearlink (RTD/CLINK origination) port. All ESCON ports continue to be configurable as host or Nearlink ports in pairs on a per CIP basis.

- As shown in [Figure G-1](#) through [Figure G-4](#), the ports are shown with their channel interface identifiers where all ports are enabled. These channel interface identifiers are the values that are required for the CHANIF values that you code for the CONFIG utility. Each value is two characters in length and has a value from 0A to 10. The first digit is the VTSS cluster ID (valid values are 0 or 1). The second digit is the group or adapter ID (valid values are A to P).

Each FICON port can attach to two RTDs, or two CLINKs, or an RTD/CLINK combination using a FICON director or supported switch (in FICON mode). Note that, as shown in these figures, **for RTDs only**, each FICON port has two CHANIF values only if the port is connected to a FICON director which is then connected to two RTDs. Nearlink RTD connections that are paired through a FICON switch or director on the same port dynamically alternate between both RTDs for atomic operations such as mount, migrate VTV, recall VTV, and so forth.

- **Each ICE card** contains two pairs of ESCON ports. Each pair is controlled by its own Channel Interface Processor (CIP). Each CIP switches between the two ports, so that **only one port** can transfer data at a time, which emulates a FICON port attached to a director attached to RTDs.
- Each host FICON channel supports 64 logical paths (times 16 logical units). However, in HCD:
 - From a single MVS host, you can only define 8 channels (CHPIDs) running to a single control unit (single VSM4).
 - You use the CNTLUNIT statement to define each VSM4 as 16 3490 control unit images.
 - You use the IODEVICE statement to define the 16 VTDS that are associated with each 3490 control unit image.
- For a VSM4, each ESCON CIP or FICON FIP can operate with only one of two modes, which is set at the VTSS LOP:

- **Host Mode.** In Host Mode, ports can connect to the host CPU channels, including using Director(s) or channel extenders. A port in Host Mode can also serve as a CLINK terminator.

Also note that for ESCON ports, you can have two physical paths from the same LPAR to the same CIP, as long as the two physical paths address different (not overlapping) logical control units. For example, a single host LPAR can address logical control units 0-7 on one CIP port, and 8-F on the other CIP port of the same CIP.

- **Nearlink Mode.** In Nearlink Mode, ports can connect to an RTD. A port in Nearlink Mode can also serve as a CLINK originator.
- **For clustering,** you need an originator port in Nearlink mode on one VTSS connected through a CLINK to a terminator port in Host mode on the other VTSS.

Caution: In bi-directional clustering, each CLINK must be attached to the same Storage Cluster on each VTSS, which is a requirement.

Failure to configure in this manner can produce Replicate, Channel, and Communication errors. For more information and examples, refer to the *ELS Disaster Recovery and Offsite Data Management Guide*.

FICON/ESCON Best Practices

For both FICON and ESCON, Best Practices for optimizing port operations are shown in [Table G-2](#).

Table G-2 Optimizing VSM4 FICON/ESCON Port Operations

Configuration - Two ESCON Ports on a CIP (ICE) or FICON port attached to a FICON Director (VCF)	Best Practices
Two CLINKs	Attach a maximum of 2 because each port allows two active operations. Note, however, that these operations share the bandwidth of the port.
CLINK and RTD	An advantage if you attach one CLINK originator/one RTD per director, because both can be active.

Table G-2 (Cont.) Optimizing VSM4 FICON/ESCON Port Operations

Configuration - Two ESCON Ports on a CIP (ICE) or FICON port attached to a FICON Director (VCF)	Best Practices
Two RTDs	<p>An advantage for the following:</p> <ul style="list-style-type: none"> ■ Optimize use of local and remote RTDs. During busy shifts, use only local RTDs on the FIP. During quiet periods, switch to remote RTDs for deep archive and DR work. Because you can have two active devices, you can also simultaneously run one local and one remote RTD. Note, however, that these operations share the bandwidth of the port. ■ Optimize use of different drive technologies. As described in the previous bullet, use a T9840 as a local RTD, then switch to a T9940 for deep archive. You can also use this feature to migrate from older drive technology (such as 9490) to newer technology (such as 9840). Use Management and Storage Classes to read in data from older media, then switch to the newer technology drive to place data on new media. This technique effectively gives you greater physical connectivity to different drive technologies without incurring the overhead of full time, real time FICON connections to each drive type. Also as above, because you can have two active devices, you can also simultaneously run two RTDs with different drive technologies. Note, however, that these operations share the bandwidth of the port.

VSM4 FICON Front-End and Back-End Configuration Examples

For VSM4s with both FICON Front-End and Back-End connectivity, look at two examples of VCF card configurations and implementation:

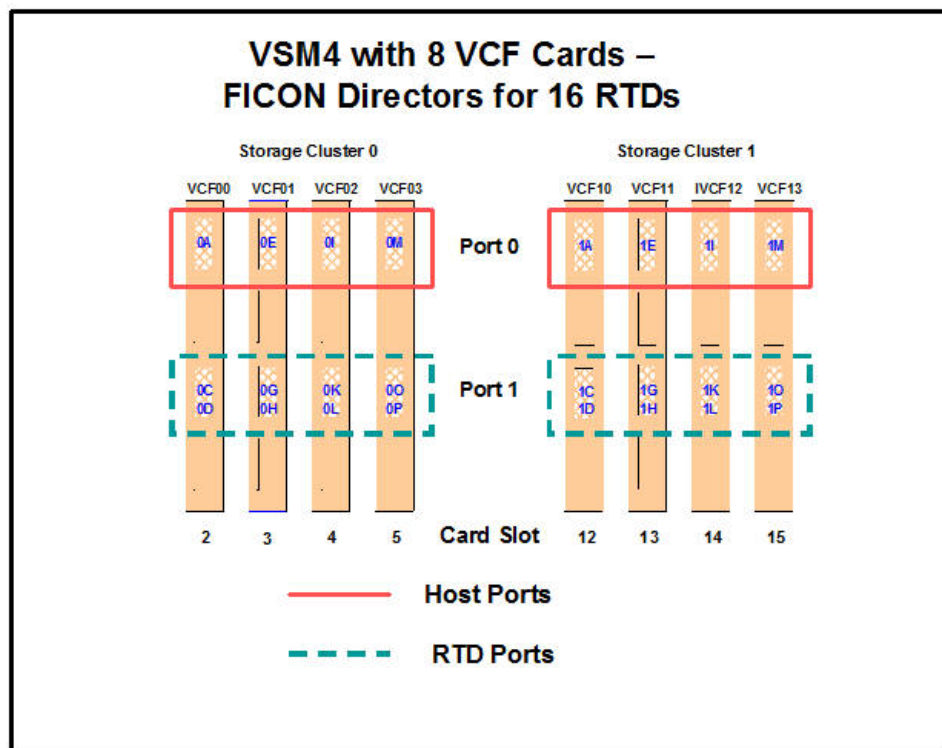
- ["VSM4 Configuration Example: 8 VCF Cards, FICON Directors, 16 RTDs"](#)
- ["VSM4 Configuration Example: 8 VCF Cards, 4 CLINKs, FICON Directors for 8 RTDs"](#)

For a VSM4 host gen example, see ["IOCP Example for Single MVS Host Connected to a VSM4 Using FICON Directors."](#)

VSM4 Configuration Example: 8 VCF Cards, FICON Directors, 16 RTDs

Figure G-5 shows CONFIG channel interface identifiers for a VSM4 with 8 VCF cards. In this configuration, you have allocated 8 ports to RTDs and 8 ports to host connections. The RTD ports are all connected to FICON directors, each of which is attached to RTDs, so the CHANIF identifiers for both RTDs are shown on each port. This allows Back-End connection to 16 RTDs, although, as with ESCON, only one RTD per port/Director can be active at a time.

Figure G-5 VSM4 with 8 VCF Cards, FICON Directors for 16 RTDs

**CONFIG Example for VSM4 FICON with 8 VCF Cards, FICON Directors, 16 RTDs**

The following example shows CONFIG JCL to define the VSM4 configuration in Figure G-5.

```
//CREATECF EXEC PGM=SLUADMIN, PARM='MIXED'
//STEPLIB DD DSN=hlq.SEALINK, DISP=SHR
//SLSCNTL DD DSN=hlq.DBASEPRM, DISP=SHR
//SLSCNTL2 DD DSN=hlq.DBASESEC, DISP=SHR
//SLSSTBY DD DSN=hlq.DBASETBY, DISP=SHR
//SLSPRINT DD SYSOUT=*
//SLSIN DD *
CONFIG
  GLOBAL MAXVTV=32000 MVCFREE=40
  RECLAIM THRESHLD=70 MAXMVC=40 START=35
  VTSS NAME=VSM401 LOW=70 HIGH=80 MAXMIG=8 RETAIN=5
  RTD NAME=VSM42A00 DEVNO=2A00 CHANIF=0C
  RTD NAME=VSM42A01 DEVNO=2A01 CHANIF=0D
  RTD NAME=VSM42A02 DEVNO=2A02 CHANIF=0G
  RTD NAME=VSM42A03 DEVNO=2A03 CHANIF=0H
  RTD NAME=VSM42A04 DEVNO=2A04 CHANIF=0K
  RTD NAME=VSM42A05 DEVNO=2A05 CHANIF=0L
  RTD NAME=VSM42A06 DEVNO=2A06 CHANIF=0O
  RTD NAME=VSM42A07 DEVNO=2A07 CHANIF=0P
  RTD NAME=VSM42A08 DEVNO=2A08 CHANIF=1C
  RTD NAME=VSM42A09 DEVNO=2A09 CHANIF=1D
  RTD NAME=VSM42A0A DEVNO=2A0A CHANIF=1G
  RTD NAME=VSM42A0B DEVNO=2A0B CHANIF=1H
  RTD NAME=VSM42A0C DEVNO=2A0C CHANIF=1K
  RTD NAME=VSM42A0D DEVNO=2A0D CHANIF=1L
```

```

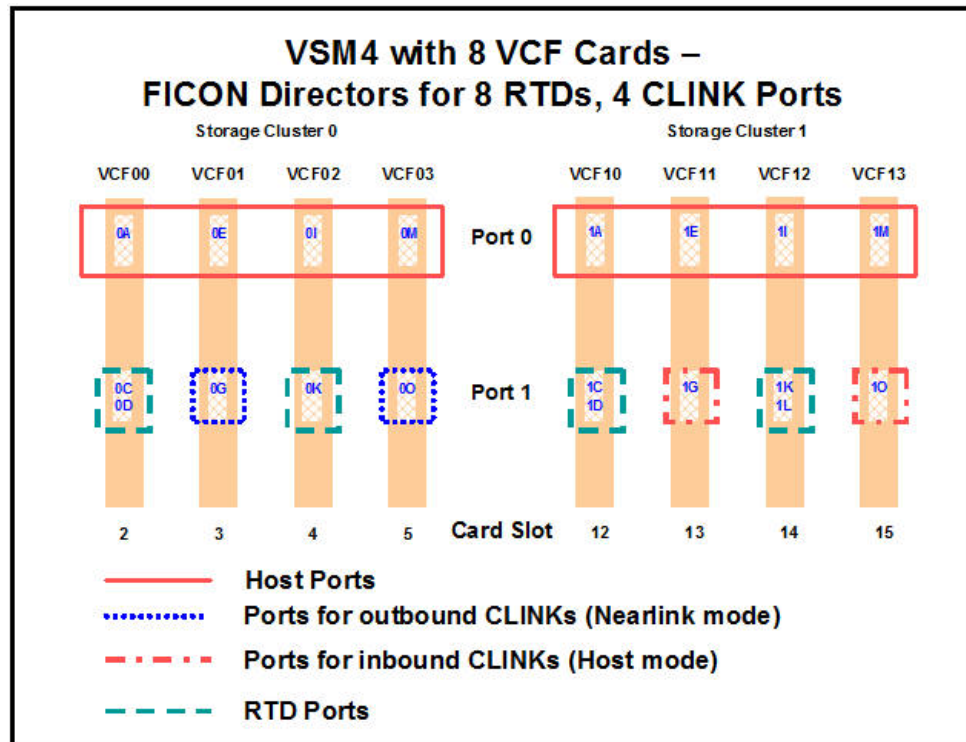
RTD  NAME=VSM42A0E  DEVNO=2A0E  CHANIF=10
RTD  NAME=VSM42A0F  DEVNO=2A0F  CHANIF=1P
VTD  LOW=9900  HIGH=99FF
    
```

VSM4 Configuration Example: 8 VCF Cards, 4 CLINKs, FICON Directors for 8 RTDs

Figure G-6 shows CONFIG channel interface identifiers for a VSM4 with 8 VCF cards. In this configuration, you have allocated:

- 8 Host ports.
- 4 ports for RTDs. The RTD ports are all connected to FICON directors, each of which is attached to RTDs, so the CHANIF identifiers for both RTDs are shown on each port. This allows Back-End connection to 8 RTDs. As with ESCON, only one RTD per port/Director can be active at a time.
- 4 ports for CLINK connections to form a Bi-Directional VTSS Cluster, and 8 ports to host connections. To form the clustered VTSS, you have two VSM4s (VSMPR1 and VSMPR2) configured identically as shown in Figure G-6. Bi-Directional Clustering requires pairs of Uni-Directional CLINKs with the FIPs configured so that the data flows in opposite directions on the CLINKs. To make that happen, make 0G and 1G the sending (Nearlink Mode) ports on both VTSSs and 0O and 1O the receiving (Host Mode) on both VTSSs.

Figure G-6 VSM4 with 8 VCF Cards, 8 Host Ports, FICON Directors for 8 RTDs, 4 CLINK Ports

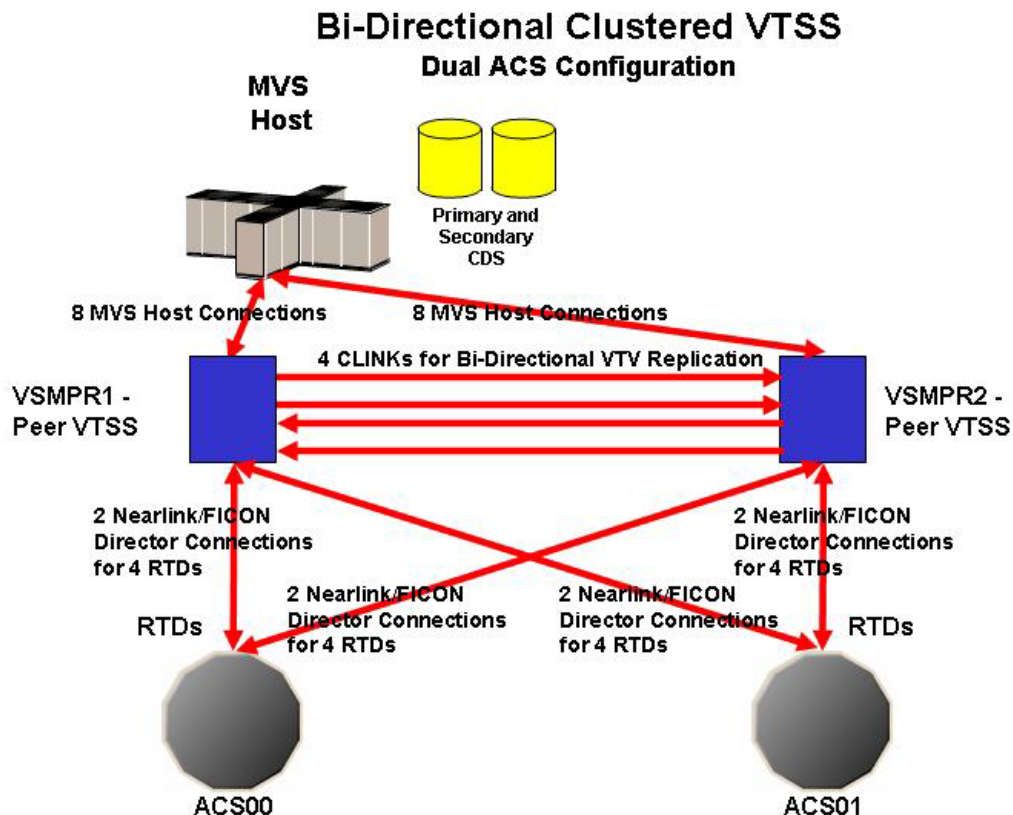


CONFIG Example for Bi-Directional Clustered VSM4 FICON Back-End

Figure G-7 shows an example of a Bi-Directional Cluster of two VSM4s (VSMPR1 and VSMPR2) with identical VCF card configurations shown in Figure G-6.

Caution: Bi-Directional Clustering requires VTCS 6.1. You cannot configure a Bi-Directional Cluster at releases lower than VTCS 6.1.

Figure G-7 Dual ACS Bi-Directional Clustered VTSS Configuration



The following example shows CONFIG JCL to define a Bi-Directional Cluster of two VSMs (VSMPR1 and VSMPR2) as shown in Figure G-7. Note that:

- The CLUSTER statement defines the Cluster as consisting of VSMPR1 and VSMPR2.
- There are CLINK statements using the sending (Nearlink Mode) ports of **both** VTSSs to enable the Cluster as Bi-Directional. The Nearlink ports are 0G and 1G on both VTSSs.

```
//CREATECF EXEC PGM=SLUADMIN, PARM='MIXED'
//STEPLIB DD DSN=hlq.SEALINK, DISP=SHR
//SLSCNTL DD DSN=hlq.DBASEPRM, DISP=SHR
//SLSCNTL2 DD DSN=hlq.DBASESEC, DISP=SHR
//SLSSTBY DD DSN=hlq.DBASETBY, DISP=SHR
//SLSPRINT DD SYSOUT=*
//SLSIN DD *
CONFIG RESET CDSLEVEL(V61ABOVE)
GLOBAL MAXVTV=32000 MVCFREE=40
RECLAIM THRESHLD=70 MAXMVC=40 START=35
VTSS NAME=VSMPR1 LOW=70 HIGH=80 MAXMIG=8 MINMIG=4 RETAIN=5
RTD NAME=PR11A00 DEVNO=1A00 CHANIF=0C
RTD NAME=PR11A01 DEVNO=1A01 CHANIF=0D
RTD NAME=PR11A02 DEVNO=1A02 CHANIF=0K
```

```


RTD  NAME=PR11A03  DEVNO=1A03  CHANIF=0L
RTD  NAME=PR12A08  DEVNO=2A08  CHANIF=1C
RTD  NAME=PR12A09  DEVNO=2A09  CHANIF=1D
RTD  NAME=PR12A0A  DEVNO=2A0A  CHANIF=1K
RTD  NAME=PR12A0B  DEVNO=2A0B  CHANIF=1L
VTD  LOW=9900  HIGH=99FF
VTSS NAME=VSMR2  LOW=70  HIGH=80  MAXMIG=8  MINMIG=4  RETAIN=5
RTD  NAME=PR23A00  DEVNO=3A00  CHANIF=0C
RTD  NAME=PR23A01  DEVNO=3A01  CHANIF=0D
RTD  NAME=PR23A02  DEVNO=3A02  CHANIF=0K
RTD  NAME=PR23A03  DEVNO=3A03  CHANIF=0L
RTD  NAME=PR24A08  DEVNO=4A08  CHANIF=1C
RTD  NAME=PR24A09  DEVNO=4A09  CHANIF=1D
RTD  NAME=PR24A0A  DEVNO=4A0A  CHANIF=1K
RTD  NAME=PR24A0B  DEVNO=4A0B  CHANIF=1L
VTD  LOW=9900  HIGH=99FF
CLUSTER NAME=CLUSTER1  VTSSs (VSMR1, VSMR2)
CLINK VTSS=VSMR1  CHANIF=0G
CLINK VTSS=VSMR1  CHANIF=1G
CLINK VTSS=VSMR2  CHANIF=0G
CLINK VTSS=VSMR2  CHANIF=1G

```

IOCP Example for Single MVS Host Connected to a VSM4 Using FICON Directors

Figure G–8 shows a configuration diagram for a single MVS host connected to a VSM4 through FICON Directors, and the example that follows shows IOCP statements for this configuration. Note that:

- From MVSA, you define 8 CHPIDs, with each path switched in the FICON Director, for a total of 8 channels running to the VSM4.
- You code 16 CNTLUNIT statements to define the VSM4 as 16 3490 images.
- You code IODEVICE statement to define the 16 VTDs that are associated with each 3490 image.
- If ESCON and FICON channels are configured to the same logical control unit, MVS issues message CBDG489I, which indicates that mixing ESCON and FICON channel paths on a logical control unit should be used only for the migration from ESCON to native FICON but should not be used permanently. This is a warning message only, and does not indicate an error.

Figure G-8 Configuration Diagram: Single MVS Host Connected to a VSM4 Using FICON Directors

```

ESCD4C CHPID PATH=(20,70),TYPE=CNC,SWITCH=4C
ESCD4D CHPID PATH=(21,71),TYPE=CNC,SWITCH=4D
ESCD4E CHPID PATH=(30,80),TYPE=CNC,SWITCH=4E
ESCD4F CHPID PATH=(31,81),TYPE=CNC,SWITCH=4F

```

```

CU1 CNTLUNIT CUNUMBR=001,
    PATH=(20,21,30,31,70,71,80,81),
    LINK=(D0,D4,D0,D4,D4,D0,D4,D0),
    UNIT=3490,CUADD=0,
    UNITADD=((00,16))

```

```

STRING1 IODEVICE ADDRESS=(0500,16),
    CUNUMBR=(001),
    UNIT=3490,
    UNITADD=00,STADET=Y

```

```

CU2 CNTLUNIT CUNUMBR=002,
    PATH=(20,21,30,31,70,71,80,81),
    LINK=(D0,D4,D0,D4,D4,D0,D4,D0),
    UNIT=3490,CUADD=1,
    UNITADD=((00,16))

```

```

STRING2 IODEVICE ADDRESS=(0510,16),
    CUNUMBR=(002),
    UNIT=3490,
    UNITADD=00,STADET=Y

```

```

.
.
.

```

```

CU15 CNTLUNIT CUNUMBR=015,
    PATH=(20,21,30,31,70,71,80,81),
    LINK=(D0,D4,D0,D4,D4,D0,D4,D0),
    UNIT=3490,CUADD=E,
    UNITADD=((00,16))

```

```

STRING15 IODEVICE ADDRESS=(05E0,16),
          CUNUMBER=(015),
          UNIT=3490,
          UNITADD=00,STADET=Y

CU16 CNTLUNIT CUNUMBR=016,
      PATH=(20,21,30,31,70,71,80,81),
      LINK=(D0,D4,D0,D4,D4,D0,D4,D0),
      UNIT=3490,CUADD=F,
      UNITADD=((00,16))

STRING16 IODEVICE ADDRESS=(05F0,16),
          CUNUMBER=(016),
          UNIT=3490,
          UNITADD=00,STADET=Y

```

Tip: Unlike ESCON, FICON supports multiple active I/Os per channel. If the number of active VTDs is less than the number of channels configured to the VTSS, the I/Os to those VTDs may not be evenly spread across all the channels. As the number of active VTDs increases to be greater than the number of channels configured to the VTSS, the channel subsystem will spread the I/Os across all the channels. If it is desired to spread the I/Os across all of the channels even when only a few VTDs are active, it is necessary to use the preferred path feature to force the channel subsystem to spread the I/Os across the channels. The preferred path feature is specified using the PATH= parameter on the IODEVICE statement. When you specify preferred path on the IODEVICE statement, the channel subsystem always tries the preferred path first. If it is busy or unavailable, the channel subsystem next tries the channel path following the preferred path in the rotation order, and so on.

The following example shows IODEVICE statements for STRING1 **without** using preferred pathing.

```

STRING1 IODEVICE ADDRESS=(0500,16),
          CUNUMBER=(001),
          UNIT=3490,
          UNITADD=00,STADET=Y

```

The following example shows IODEVICE statements for STRING1 using preferred pathing. If you are using preferred pathing, you need to use these kind of IODEVICE statements for all paths, such as STRING2 through STRING16 shown in the previous example.

```

STRING10 IODEVICE ADDRESS=(0500,2),
          CUNUMBER=(001),
          UNIT=3490,
          UNITADD=00,STADET=Y,
          PATH=20

STRING12 IODEVICE ADDRESS=(0502,2),
          CUNUMBER=(001),
          UNIT=3490,
          UNITADD=00,STADET=Y,
          PATH=21

STRING14 IODEVICE ADDRESS=(0504,2),
          CUNUMBER=(001),
          UNIT=3490,
          UNITADD=00,STADET=Y,

```

```
PATH=30

STRING16 IODEVICE ADDRESS=(0506,2),
        CUNUMBER=(001),
        UNIT=3490,
        UNITADD=00,STADET=Y,
        PATH=31

STRING18 IODEVICE ADDRESS=(0508,2),
        CUNUMBER=(001),
        UNIT=3490,
        UNITADD=00,STADET=Y,
        PATH=70

STRING1A IODEVICE ADDRESS=(050A,2),
        CUNUMBER=(001),
        UNIT=3490,
        UNITADD=00,STADET=Y,
        PATH=71

STRING1C IODEVICE ADDRESS=(050C,2),
        CUNUMBER=(001),
        UNIT=3490,
        UNITADD=00,STADET=Y,
        PATH=80

STRING1E IODEVICE ADDRESS=(050E,2),
        CUNUMBER=(001),
        UNIT=3490,
        UNITADD=00,STADET=Y,
        PATH=81
```

VSM5 FICON Configuration

The VSM5, provides greater capacity and throughput than the VSM4, while retaining its advantages over the VSM3. [Table H-1](#) summarizes the VSM5 features.

Table H-1 VSM5 Features

Feature	Description
Host/Nearlink Interfaces	Up to 16 (FICON only)
RTDs supported	Up to 32 through FICON directors (in 3490-emulation mode only), can be a mixture of the following: 9840B, 9840C, 9940B, T10000.
LSMs supported	9740, 9360, 4410, 9310, SL8500 , SL3000
Maximum VTDs per VTSS	256
Maximum VTVs per VTSS	300,000

VSM5 FICON VCF Card Options - Maximum 16 RTDs

VSM5 is available only with VCF (FICON) cards in the following configurations for a maximum of 16 RTDs:

- [Figure H-1](#) shows a VSM5 with 8 VCF cards.
- [Figure H-2](#) shows a VSM5 with 6 VCF cards, 2 empty card slots.
- [Figure H-3](#) shows a VSM5 with 4 VCF cards, 4 empty card slots.

Figure H-1 VSM5 with 8 VCF Cards - Max 16 RTDs

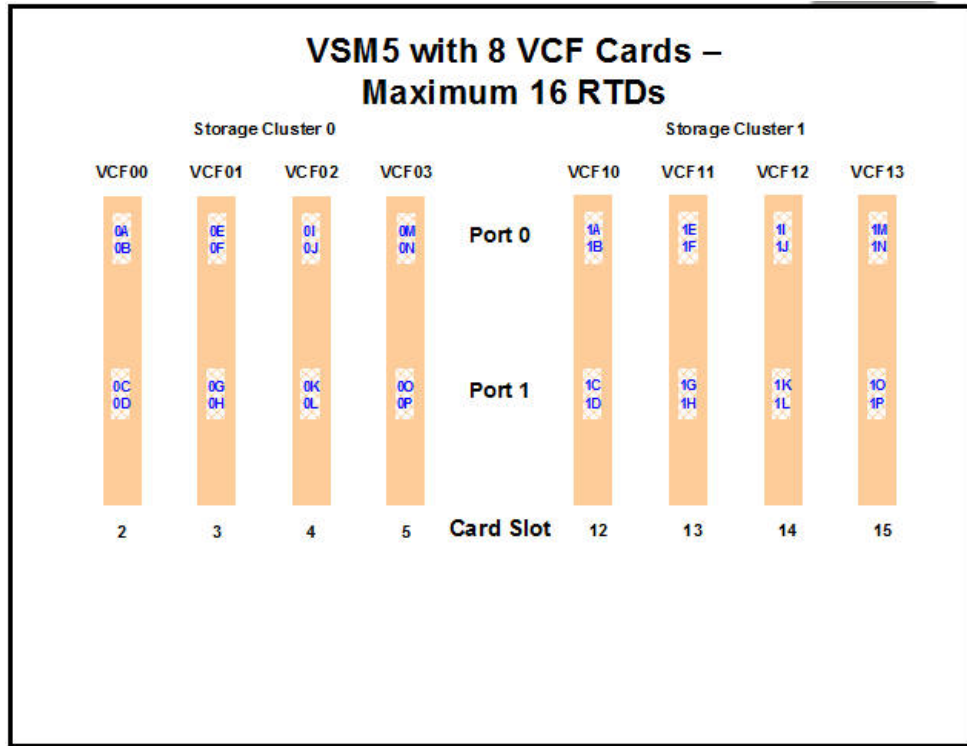


Figure H-2 VSM5 with 6 VCF Cards, 2 Empty Card Slots

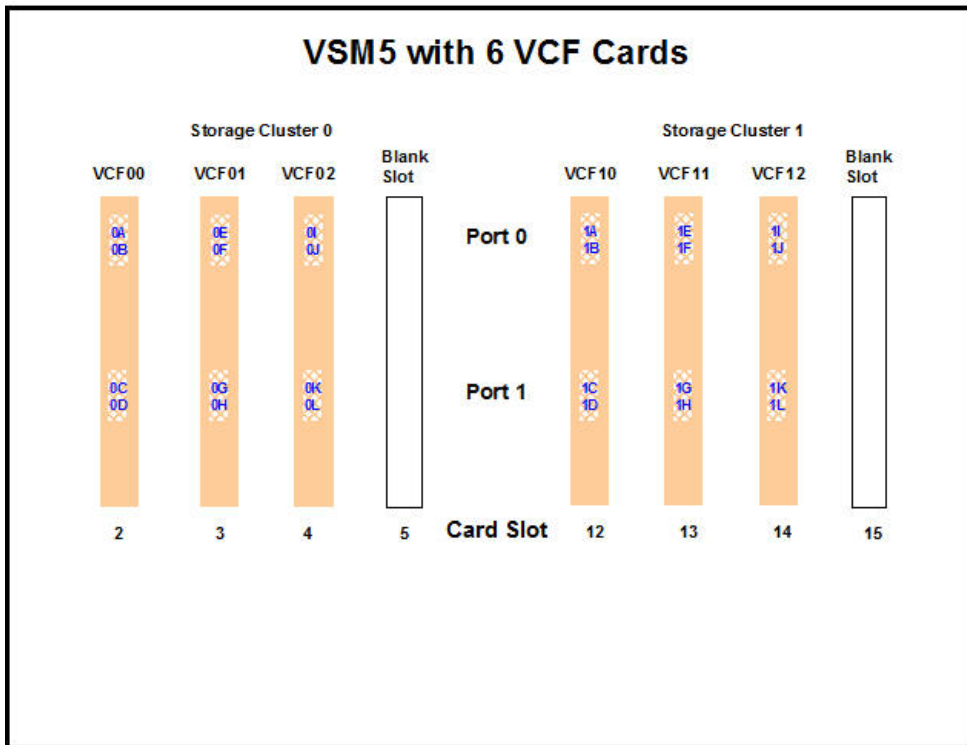



Figure H-3 VSM5 with 4 VCF Cards, 4 Empty Card Slots




Note: In Figure H-1 through Figure H-3, the VCF cards must go in:

- All slots in an eight-VCF card configuration.
- Slots 2, 3, 4, 13, 14, and 15 in a six-VCF card configuration.
- Slots 2, 4, 14, and 15 in a four-VCF card configuration.

VSM5 FICON VCF Card Options - Maximum 32 RTDs

VSM5 is available **only** with 8 VCF (FICON) cards in the configuration for a maximum of 32 RTDs shown in Figure H-4. For more information on device addressing, see "RTD/CLINK Addresses - Maximum 32 RTDs."

Figure H-4 VSM5 with 8 VCF cards - Max 32 RTDs



FICON Port Processing

FICON ports processing is as follows:

- FICON ports are controlled by a FICON Interface processor (FIP) and there can be only a total of 14 Nearlink FIPs (for maximum of 16 RTDs supported) or 28 Nearlink FIPs (for maximum of 32 RTDs supported).
- For a VSM5, each FIP can operate with only one of two modes, which is set at the VTSS DOP:
 - **Host Mode.** In Host Mode, ports can connect to the host CPU channels, including through Director(s) or channel extenders. A port in Host Mode can also serve as a CLINK terminator.
 - **Nearlink Mode.** In Nearlink Mode, ports can connect to an RTD. A port in Nearlink Mode can also serve as a CLINK originator.
 - **For clustering,** you need an originator port in Nearlink mode on one VTSS connected through a CLINK to a terminator port in Host mode on the other VTSS.
- In Figure H-1 through Figure H-4, the ports are shown with their channel interface identifiers where all ports are enabled. For more information on device addressing, see "RTD/CLINK Addresses - Maximum 32 RTDs."

Each FICON port can attach to up to 4 RTDs, or up to 4 CLINKs, or up to 4 RTD/CLINK combinations through a FICON director or supported switch (in FICON mode). Note that, as shown in these figures, each FICON port has multiple device addresses **only if** the port is connected to a FICON director which is then connected to multiple devices.

Note: With microcode level D02.06.00.00 or higher, multiple Nearlink device connections through a FICON switch or Director on the same port now allow:

- **Up to a total of 16 simultaneous NearLink I/O transfers**, which can be spread across multiple targets on as many as 14 NearLink ports.
 - **Up to a total of 2 simultaneous NearLink I/O transfers are allowed per port.**
-
-
- Each host FICON channel supports 64 logical paths (times 16 logical units). However, in HCD:
 - From a single MVS host, you can only define 8 channels (CHPIDs) running to a single control unit (single VSM5).
 - You use the CNTLUNIT statement to define each VSM5 as 16 3490 control unit images.
 - You use the IODEVICE statement to define the 16 VTDs that are associated with each 3490 control unit image.

Caution: In bi-directional clustering, each CLINK must be attached to the same Storage Cluster on each VTSS, which is a requirement. Failure to configure in this manner can produce Replicate, Channel, and Communication errors. For more information and examples, refer to the *ELS Disaster Recovery and Offsite Data Management Guide*.

FICON Port Operations Best Practices

For FICON, Best Practices for optimizing port operations are shown in [Table H-2](#).

Table H-2 *Optimizing VSM5 FICON Port Operations*

Configuration - FICON port attached to a FICON Director	Best Practices
Multiple CLINKs (up to 4)	Attach a maximum of 2 because each port allows two active operations. Note, however, that these operations share the bandwidth of the port.
CLINK and RTD combinations	An advantage if you attach one CLINK originator /one RTD per director, because both can be active.

Table H-2 (Cont.) Optimizing VSM5 FICON Port Operations

Configuration - FICON port attached to a FICON Director	Best Practices
Up to 4 RTDs	<p>An advantage for the following:</p> <ul style="list-style-type: none"> ■ Optimize use of local and remote RTDs. During busy shifts, use only local RTDs on the FIP. During quiet periods, switch to remote RTDs for deep archive and DR work. Because you can have two active devices, you can also simultaneously run one local and one remote RTD. Note, however, that these operations share the bandwidth of the port. ■ Optimize use of different drive technologies. As described in the previous bullet, use a T9840 as a local RTD, then switch to a T9940 for deep archive. You can also use this feature to migrate from older drive technology (such as 9490) to newer technology (such as 9840). Use Management and Storage Classes to read in data from older media, then switch to the newer technology drive to place data on new media. This technique effectively gives you greater physical connectivity to different drive technologies without incurring the overhead of full time, real time FICON connections to each drive type. Also as above, because you can have two active devices, you can also simultaneously run two RTDs with different drive technologies. Note, however, that these operations share the bandwidth of the port.

VSM5 FICON Front-End and Back-End Configuration Examples

For VSM5s, look at an examples of VCF card configurations and implementation in ["VSM5 with 8 VCF Cards, FICON Directors for 32 RTDs."](#)

For a VSM5 host gen example, see ["IOCP Example for Single MVS Host Connected to a VSM5 Using FICON Directors."](#)

Implementing Support for Maximum of 32 RTDs

1. Ensure that your system has the Maximum 32 RTDs requirements described in [Table H-3](#).

Table H-3 32 RTDs/Stacked Migrates Support Requirements for VTCS/NCS 7.0

32 RTDs/Stacked Migrates Support Requirement	VSM4/VSM5 Microcode Level	CDS Level
FICON RTDs and FICON ports for the CLINKs	D02.05.00.00 or higher	"F" or higher

2. Use CONFIG GLOBAL to enable support for maximum of 32 RTDs.

```
CONFIG GLOBAL MAXRTDS=32
```

Note: Enabling support for a maximum of 32 RTDs does not require CONFIG RESET. However, regressing from 32 RTDs supported to 16 RTDs supported does require CONFIG RESET.

3. Update your CONFIG RTD and CONFIG CLINK statements as required. For more information, see:
 - ["Maximum RTDs per VTSS"](#)
 - ["VSM5 Configuration Example: 8 VCF Cards, FICON Directors, 32 RTDs"](#)

Note: The CONFIG utility RTD statement defines the RTDs connected to the VTSS. Specifically, the CONFIG RTD CHANIF parameter specifies the channel interface on the VTSS that communicates with the RTD.

Similarly, the The CONFIG utility CLINK statement defines the channel interface for a CLINK originator using the CONFIG CLINK CHANIF parameter.

Code values for the CHANIF parameter as follows:

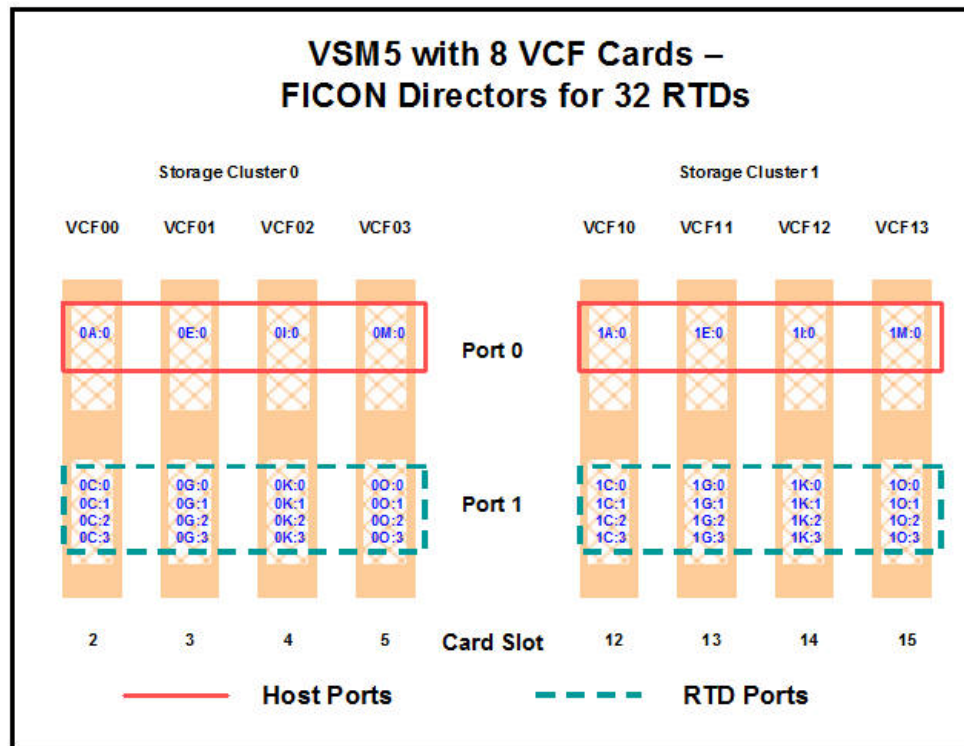
- Regardless of whether the Maximum 32 RTDs feature is enabled, if you do not have a total of greater than 16 (RTDs, CLINK originators, or a combination of RTDs and CLINK originators) on that VTSS, you can use the "old" addressing scheme on the CHANIF parameters.
 - If, however, the Maximum 32 RTDs feature is enabled and you have total of greater than 16 (RTDs, CLINK originators, or a combination of RTDs and CLINK originators) on that VTSS, you must use the "new" addressing scheme on the corresponding CHANIF parameters.
-

4. Use the VSM5 DOP to reenter your RTD device addresses. See ["VSM5 DOP Panels for Maximum 32 RTDs."](#)

VSM5 Configuration Example: 8 VCF Cards, FICON Directors, 32 RTDs

[Figure H-5](#) shows CONFIG channel interface identifiers for a VSM5 with 8 VCF cards and the Maximum 32 RTDs feature enabled. In this configuration, you have allocated 8 ports to RTDs and 8 ports to host connections. The RTD ports are all connected to FICON directors, each of which is attached to 4 RTDs, so the CHANIF identifiers for all 4 RTDs are shown on each port. This allows Back-End connection to 32 RTDs, although only one RTD per port/Director can be active at a time.

Figure H-5 VSM5 with 8 VCF Cards, FICON Directors for 32 RTDs



CONFIG Example for VSM5 FICON with 8 VCF Cards, FICON Directors, 32 RTDs

The following example shows CONFIG JCL to define the VSM5 configuration shown in Figure H-5.

```
//CREATECF EXEC PGM=SLUADMIN, PARM='MIXED'
//STEPLIB DD DSN=hlq.SEALINK, DISP=SHR
//SLSCNTL DD DSN=hlq.DBASEPRM, DISP=SHR
//SLSCNTL2 DD DSN=hlq.DBASESEC, DISP=SHR
//SLSSTBY DD DSN=hlq.DBASESTBY, DISP=SHR
//SLSPRINT DD SYSOUT=*
//SLSIN DD *
CONFIG
GLOBAL MAXVTV=32000 MCVFREE=40 VTVATTR=SCRATCH RECALWER=YES LOCKSTR=VTCS_LOCKS
REPLICAT=ALWAYS VTVPAGE=LARGE SYNCHREP=YES MAXRTDS=32
RECLAIM THRESHLD=70 MAXMVC=40 START=35
VTSS NAME=VSM501 LOW=70 HIGH=80 MAXMIG=8 RETAIN=5
RTD NAME=VSM52A00 DEVNO=2A00 CHANIF=0C:0
RTD NAME=VSM52A01 DEVNO=2A01 CHANIF=0C:1
RTD NAME=VSM52A02 DEVNO=2A02 CHANIF=0C:2
RTD NAME=VSM52A03 DEVNO=2A03 CHANIF=0C:3
RTD NAME=VSM52A04 DEVNO=2A04 CHANIF=0G:0
RTD NAME=VSM52A05 DEVNO=2A05 CHANIF=0G:1
RTD NAME=VSM52A06 DEVNO=2A06 CHANIF=0G:2
RTD NAME=VSM52A07 DEVNO=2A07 CHANIF=0G:3
RTD NAME=VSM52A08 DEVNO=2A08 CHANIF=0K:0
RTD NAME=VSM52A09 DEVNO=2A09 CHANIF=0K:1
RTD NAME=VSM52A0A DEVNO=2A0A CHANIF=0K:2
RTD NAME=VSM52A0B DEVNO=2A0B CHANIF=0K:3
RTD NAME=VSM52A0C DEVNO=2A0C CHANIF=0O:0
RTD NAME=VSM52A0D DEVNO=2A0D CHANIF=0O:1
```

```
RTD  NAME=VSM52A0E  DEVNO=2A0E  CHANIF=00:2
RTD  NAME=VSM52A0F  DEVNO=2A0F  CHANIF=00:3
RTD  NAME=VSM53A00  DEVNO=3A00  CHANIF=1C:0
RTD  NAME=VSM53A01  DEVNO=3A01  CHANIF=1C:1
RTD  NAME=VSM53A02  DEVNO=3A02  CHANIF=1C:2
RTD  NAME=VSM53A03  DEVNO=3A03  CHANIF=1C:3
RTD  NAME=VSM53A04  DEVNO=3A04  CHANIF=1G:0
RTD  NAME=VSM53A05  DEVNO=3A05  CHANIF=1G:1
RTD  NAME=VSM53A06  DEVNO=3A06  CHANIF=1G:2
RTD  NAME=VSM53A07  DEVNO=3A07  CHANIF=1G:3
RTD  NAME=VSM53A08  DEVNO=3A08  CHANIF=1K:0
RTD  NAME=VSM53A09  DEVNO=3A09  CHANIF=1K:1
RTD  NAME=VSM53A0A  DEVNO=3A0A  CHANIF=1K:2
RTD  NAME=VSM53A0B  DEVNO=3A0B  CHANIF=1K:3
RTD  NAME=VSM53A0C  DEVNO=3A0C  CHANIF=10:0
RTD  NAME=VSM53A0D  DEVNO=3A0D  CHANIF=10:1
RTD  NAME=VSM53A0E  DEVNO=3A0E  CHANIF=10:2
RTD  NAME=VSM53A0F  DEVNO=3A0F  CHANIF=10:3
VTD  LOW=9900  HIGH=99FF
```

VSM5 DOP Panels for Maximum 32 RTDs

The following VSM5 DOP panels are shown here.

Channel Configuration Status Screen To access the Channel Configuration Status screen, click the active Channel Status text field on the Configuration/Status Menu screen. See [Figure H-6](#).

Figure H-6 Channel Configuration Status Screen

STORAGETEK™ VSM - Virtual Storage Manager

Status: Full Box IML Complete | IP: 129.80.70.9 | S/N: 0567-00200047 | Master ISP: 0

Channel Configuration Status

Card	Name	CL	Lk	Gr	En	Type	RTD Port	ID
VCF00		0	0	A	Y	HOST		
		0	0	B	N			
		0	1	C	Y	HOST		
		0	1	D	N			
VCF01		0	0	E	Y	NEARLINK	00	22
							FF	FF
							FF	FF
							FF	FF
VCF02		0	0	F	N			
		0	1	G	Y	HOST		
		0	1	H	N			
		0	0	I	Y	HOST		
VCF03		0	0	J	N			
		0	1	K	Y	HOST		
		0	1	L	N			
VCF10		0	0	M	Y	HOST		
		0	0	N	N			
		0	1	O	Y	HOST		
VCF11		0	1	P	N			
		1	0	A	Y	HOST		
		1	0	B	N			
		1	1	C	Y	HOST		
VCF12		1	1	D	N			
		1	0	E	Y	NEARLINK	61	20
							FF	FF
							FF	FF
VCF13							FF	FF
		1	0	F	N			
		1	1	G	Y	HOST		
		1	1	H	N			
VCF13		1	0	I	Y	HOST		
		1	0	J	N			
		1	1	K	Y	HOST		
		1	1	L	N			
VCF13		1	0	M	Y	NEARLINK	00	00
							FF	FF
							FF	FF
							FF	FF
VCF13		1	0	N	N			
		1	1	O	Y	HOST		
		1	1	P	N			

Buttons: Exit, Main, Help, FSC/DCC, hic_stat

Channel Configuration and RTD Path Validation Screen To access the Channel Configuration and RTD Path Validation screen, click on a VCF card shown on the Channel Configuration Status screen.

To set the configuration of a VCF card channel for host or Nearlink use, select the channel (0 or 1) and type from the pull-down lists, then click Continue to display a subscreen with the message Success, indicating the configuration change completed successfully. Click Cancel to undo changed settings and return to the Channel Configuration Status screen. See [Figure H-7](#).

Figure H-7 Channel Configuration and RTD Path Validation Screen

The screenshot displays the StorageTek VSM Op-Panel interface. At the top, the browser title is "StorageTek: VSM Op-Panel - Microsoft Internet Explorer provided by StorageTek". The main header shows "STORAGETEK™ VSM - Virtual Storage Manager".

Status Bar:

Status Full Box IML Complete	IP 129.80.70.9	S/N 0567-00200047	Master ISP 0	
--	--------------------------	-----------------------------	------------------------	--

Channel Configuration

Card:
Channel:
Name:
Cluster: 0
Link: 0
Group: E
Enable: true
Type:

RTD0 DD: AA:
RTD1 DD: AA:
RTD2 DD: AA:
RTD3 DD: AA:

RTD Path Validation

Validation Path:

Navigation:

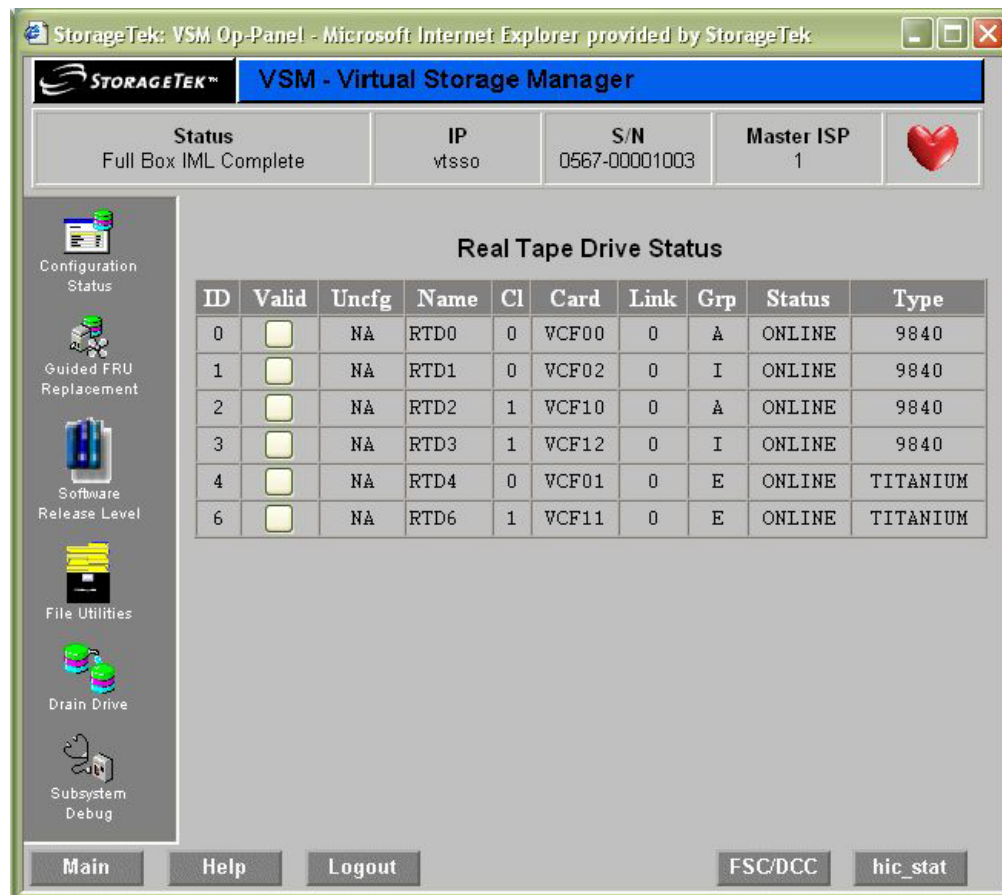
- Exit
- Configuration Status
- Guided FRU Replacement
- Software Release Level
- File Utilities
- Drain Drive
- Subsystem Debug

Footer:

Real Tape Drive Status Screen To access the Real Tape Drive Status screen, click the active Real Tape Drive Status text field on the Configuration / Status Menu screen, [Figure H-8](#). To validate a real tape drive (RTD), click the active button in the Valid column for the RTD. The VTSS support facility validates the RTD, then displays a subscreen with the message RTD n was successfully validated. See hic_stat for details.

Note: RTD configuration is preserved/restored across cold IML (EPO or CPD) but the links are reset and RTDs are offline until you vary them online with the VTCS VARY RTD ONLINE command.

Figure H-8 Real Tape Drive Status Screen



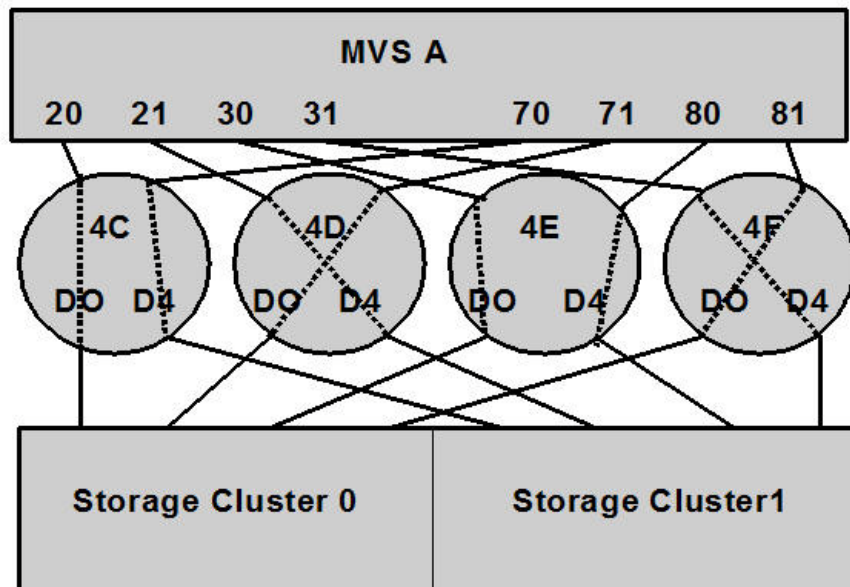
IOCP Example for Single MVS Host Connected to a VSM5 Using FICON Directors

[Figure H-9](#) shows a configuration diagram for a single MVS host connected to a VSM5 through FICON Directors, and the example that follows it shows IOCP statements for this configuration. Note that:

- From MVSA, you define 8 CHPIDs, with each path switched in the FICON Director, for a total of 8 channels running to the VSM5.
- You code 16 CNTLUNIT statements to define the VSM5 as 16 3490 images.
- You code IODEVICE statement to define the 16 VTDs that are associated with each 3490 image.

- If ESCON and FICON channels are configured to the same logical control unit, MVS issues message CBDG489I, which indicates that mixing ESCON and FICON channel paths on a logical control unit should be used only for the migration from ESCON to native FICON but should not be used permanently. This is a warning message only, and does not indicate an error.

Figure H-9 Configuration Diagram: Single MVS Host Connected to a VSM5 Using FICON Directors



```

ESCD4C CHPID PATH=(20,70),TYPE=FC,SWITCH=4C
ESCD4D CHPID PATH=(21,71),TYPE=FC,SWITCH=4D
ESCD4E CHPID PATH=(30,80),TYPE=FC,SWITCH=4E
ESCD4F CHPID PATH=(31,81),TYPE=FC,SWITCH=4F

```

```

CU1 CNTLUNIT CUNUMBR=001,
  PATH=(20,21,30,31,70,71,80,81),
  LINK=(D0,D4,D0,D4,D4,D0,D4,D0),
  UNIT=3490,CUADD=0,
  UNITADD=((00,16))

```

```

STRING1 IODEVICE ADDRESS=(0500,16),
  CUNUMBR=(001),
  UNIT=3490,
  UNITADD=00,STADET=Y

```

```

CU2 CNTLUNIT CUNUMBR=002,
  PATH=(20,21,30,31,70,71,80,81),
  LINK=(D0,D4,D0,D4,D4,D0,D4,D0),
  UNIT=3490,CUADD=1,
  UNITADD=((00,16))

```

```

STRING2 IODEVICE ADDRESS=(0510,16),
  CUNUMBR=(002),
  UNIT=3490,
  UNITADD=00,STADET=Y

```

```

.
.

```

```
CU15 CNTRLUNIT CUNUMBR=015,  
      PATH=(20,21,30,31,70,71,80,81),  
      LINK=(D0,D4,D0,D4,D4,D0,D4,D0),  
      UNIT=3490,CUADD=E,  
      UNITADD=((00,16))
```

```
STRING15 IODEVICE ADDRESS=(05E0,16),  
          CUNUMBER=(015),  
          UNIT=3490,  
          UNITADD=00,STADET=Y
```

```
CU16 CNTRLUNIT CUNUMBR=016,  
      PATH=(20,21,30,31,70,71,80,81),  
      LINK=(D0,D4,D0,D4,D4,D0,D4,D0),  
      UNIT=3490,CUADD=F,  
      UNITADD=((00,16))
```

```
STRING16 IODEVICE ADDRESS=(05F0,16),  
          CUNUMBER=(016),  
          UNIT=3490,  
          UNITADD=00,STADET=Y
```

Tip: Unlike ESCON, FICON supports multiple active I/Os per channel. If the number of active VTDs is less than the number of channels configured to the VTSS, the I/Os to those VTDs may not be evenly spread across all the channels. As the number of active VTDs increases to be greater than the number of channels configured to the VTSS, the channel subsystem will spread the I/Os across all the channels. If it is desired to spread the I/Os across all of the channels even when only a few VTDs are active, it is necessary to use the preferred path feature to force the channel subsystem to spread the I/Os across the channels. The preferred path feature is specified using the PATH= parameter on the IODEVICE statement. When you specify preferred path on the IODEVICE statement, the channel subsystem always tries the preferred path first. If it is busy or unavailable, the channel subsystem next tries the channel path following the preferred path in the rotation order, and so on.

The following example shows IODEVICE statements for STRING1 without using preferred pathing.

```
STRING1 IODEVICE ADDRESS=(0500,16),  
          CUNUMBER=(001),  
          UNIT=3490,  
          UNITADD=00,STADET=Y
```

The following example shows IODEVICE statements for STRING1 using preferred pathing. If you are using preferred pathing, you need to use these kind of IODEVICE statements for all paths, such as STRING2 through STRING16.

```
STRING10 IODEVICE ADDRESS=(0500,2),  
          CUNUMBER=(001),  
          UNIT=3490,  
          UNITADD=00,STADET=Y,  
          PATH=20
```

```
STRING12 IODEVICE ADDRESS=(0502,2),  
          CUNUMBER=(001),  
          UNIT=3490,  
          UNITADD=00,STADET=Y,
```

```
PATH=21

STRING14 IODEVICE ADDRESS=(0504,2),
        CUNUMBER=(001),
        UNIT=3490,
        UNITADD=00,STADET=Y,
        PATH=30

STRING16 IODEVICE ADDRESS=(0506,2),
        CUNUMBER=(001),
        UNIT=3490,
        UNITADD=00,STADET=Y,
        PATH=31

STRING18 IODEVICE ADDRESS=(0508,2),
        CUNUMBER=(001),
        UNIT=3490,
        UNITADD=00,STADET=Y,
        PATH=70

STRING1A IODEVICE ADDRESS=(050A,2),
        CUNUMBER=(001),
        UNIT=3490,
        UNITADD=00,STADET=Y,
        PATH=71

STRING1C IODEVICE ADDRESS=(050C,2),
        CUNUMBER=(001),
        UNIT=3490,
        UNITADD=00,STADET=Y,
        PATH=80

STRING1E IODEVICE ADDRESS=(050E,2),
        CUNUMBER=(001),
        UNIT=3490,
        UNITADD=00,STADET=Y,
        PATH=81
```

VSM5 ESCON/FICON Configurations

[Table I-1](#) summarizes the supported VCF (FICON) and ICE (ESCON) card configurations for VSM5. VSM5 ESCON support requires microcode level D02.07.00.00 or H01.07.00.00.

Table I-1 Supported Card Configurations for VSM4 FICON Front-End plus Back-End Connectivity

VCF Cards	FICON Ports	ICE Cards	ESCON Ports	Total Ports	Total Logical Paths (16 per ICE Port, 64 per VCF Port)
0	0	8	32	32	512
4	8	4	16	24	768

VSM5 ICE/VCF Card Options

VSM5 supports the following ICE/VCF card options:

- [Figure I-1](#) shows a VSM5 with 8 ICE cards.
- [Figure I-2](#) shows a VSM5 with 4 ICE cards, 4 VCF cards.

Figure I-1 Supported Card Configurations for VSM5 ESCON/FICON





Figure I-2 VSM5 with 4 ICE cards, 4 VCF cards



- [Figure I-2](#), the VCF cards must go in: Slots 3, 4, 13, and 14 in a four-VCF card configuration.
- FICON ports are controlled by a FICON Interface processor (FIP), ESCON ports are controlled by a CIP. Regardless of the card configuration, there can be only a total of 14 Nearlink FIPs and/or CIPs.

Note: With microcode level D02.06.00.00 or higher, multiple Nearlink device connections through a FICON or ESCON switch or Director on the same port now allow:

- **Up to a total of 16 simultaneous NearLink I/O transfers**, which can be spread across multiple targets on as many as 14 NearLink ports.
 - **Up to a total of 2 simultaneous NearLink I/O transfers are allowed per port.**
-
- All FICON ports can be configured as either a Host port or Nearlink (RTD/CLINK origination) port. All ESCON ports continue to be configurable as host or Nearlink ports in pairs on a per CIP basis.
 - As shown in [Figure I-1](#) and [Figure I-2](#), the ports are shown with their channel interface identifiers. These channel interface identifiers are the values that are

required for the CHANIF values that you code for the CONFIG utility. Each value is two characters in length and has a value from 0A to 1O. The first digit is the VTSS cluster ID (valid values are 0 or 1). The second digit is the group or adapter ID (valid values are A to P).

Each FICON port can attach to four RTDs, or two CLINKs, or an RTD/CLINK combination through a FICON director or supported switch (in FICON mode). Note that, as shown in these figures, **for RTDs only**, each FICON port has two CHANIF values only if the port is connected to a FICON director which is then connected to two RTDs. Nearlink RTD connections that are paired through a FICON switch or director on the same port dynamically alternate between both RTDs for atomic operations such as mount, migrate VTV, recall VTV, and so forth.

- **Each ICE card** contains two pairs of ESCON ports. Each pair is controlled by its own Channel Interface Processor (CIP). Each CIP switches between the two ports, so that **only one port** can transfer data at a time, which emulates a FICON port attached to a director attached to RTDs.
- Each host FICON channel supports 64 logical paths (times 16 logical units). However, in HCD:
 - From a single MVS host, you can only define 8 channels (CHPIDs) running to a single control unit (single VSM5).
 - You use the CNTLUNIT statement to define each VSM5 as 16 3490 control unit images.
 - You use the IODEVICE statement to define the 16 VTDs that are associated with each 3490 control unit image.
- For a VSM5, each ESCON CIP or FICON FIP can operate with only one of two modes, which is set at the VTSS DOP:
 - **Host Mode.** In Host Mode, ports can connect to the host CPU channels, including using Director(s) or channel extenders. A port in Host Mode can also serve as a CLINK terminator.

Also note that for ESCON ports, you can have two physical paths from the same LPAR to the same CIP, as long as the two physical paths address different (not overlapping) logical control units. For example, a single host LPAR can address logical control units 0-7 on one CIP port, and 8-F on the other CIP port of the same CIP.
 - **Nearlink Mode.** In Nearlink Mode, ports can connect to an RTD. A port in Nearlink Mode can also serve as a CLINK originator.

Caution: In bi-directional clustering, each CLINK must be attached to the same Storage Cluster on each VTSS, which is a requirement. Failure to configure in this manner can produce Replicate, Channel, and Communication errors.

For both FICON and ESCON, [Table I-2](#) shows Best Practices for optimizing port operations.

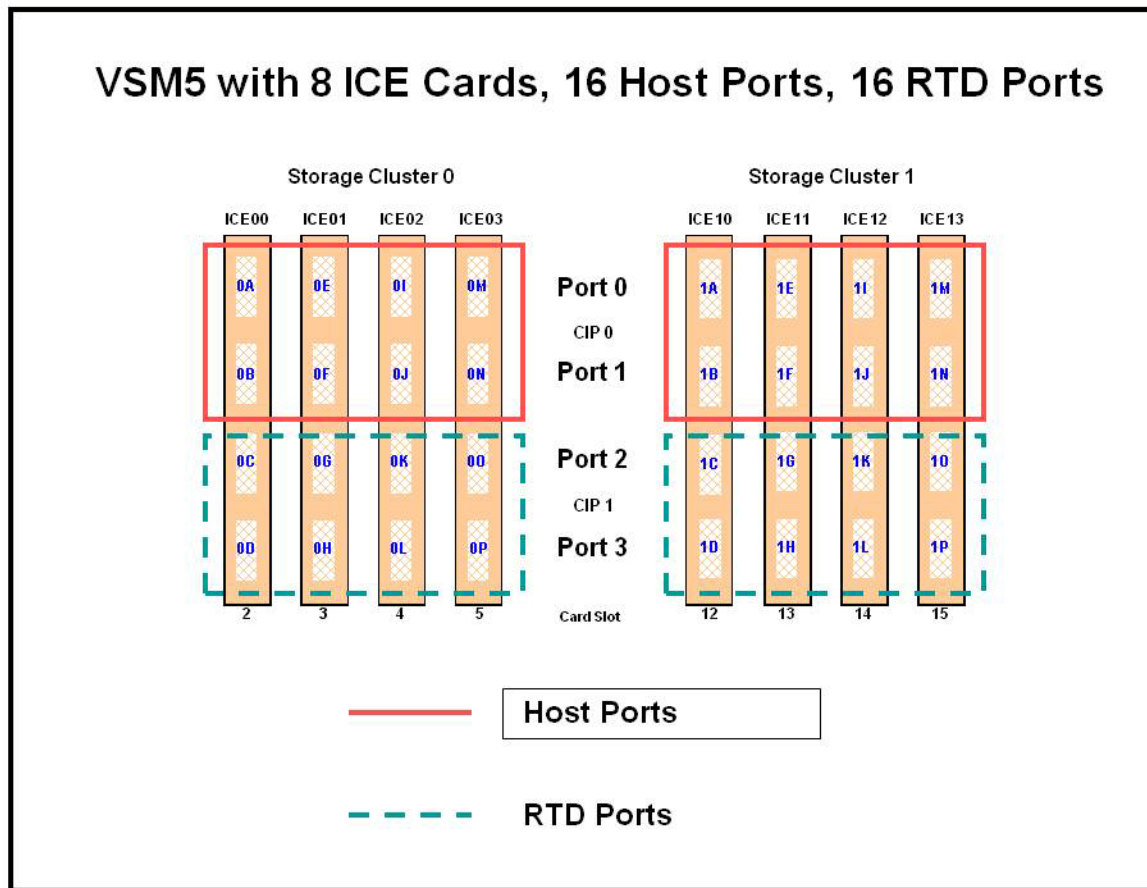
Table I-2 Optimizing VSM5 FICON/ESCON Port Operations

Configuration - Two ESCON Ports on a CIP (ICE) or FICON port attached to a FICON Director (VCF)	Best Practices
Multiple CLINKs (up to 4)	Attach a maximum of 2 because each port allows two active operations. Note, however, that these operations share the bandwidth of the port.
CLINK and RTD combinations	An advantage if you attach one CLINK originator /one RTD per director, because both can be active.
Up to 4 RTDs	<p>An advantage for the following:</p> <ul style="list-style-type: none"> ▪ Optimize use of local and remote RTDs. During busy shifts, use only local RTDs on the FIP. During quiet periods, switch to remote RTDs for deep archive and DR work. Because you can have two active devices, you can also simultaneously run one local and one remote RTD. Note, however, that these operations share the bandwidth of the port. ▪ Optimize use of different drive technologies. As described in the previous bullet, use a T9840 as a local RTD, then switch to a T9940 for deep archive. You can also use this feature to migrate from older drive technology (such as 9490) to newer technology (such as 9840). Use Management and Storage Classes to read in data from older media, then switch to the newer technology drive to place data on new media. This technique effectively gives you greater physical connectivity to different drive technologies without incurring the overhead of full time, real time FICON connections to each drive type. Also as above, because you can have two active devices, you can also simultaneously run two RTDs with different drive technologies. Note, however, that these operations share the bandwidth of the port..

VSM5 Configuration Example: 8 ICE Cards, 16 Host Ports, 16 RTD Ports

Figure I-3 shows CONFIG channel interface identifiers of 16 for hosts, 16 for RTDs for a VSM5.

Figure I-3 VSM5 with 16 Host Ports, 16 RTD Ports



CONFIG Example for VSM5 with 16 Host Ports, 16 RTD Ports

The following example shows CONFIG JCL to define the VSM5 configuration shown in Figure I-3.

```
//CREATECF EXEC PGM=SLUADMIN, PARM='MIXED'
//STEPLIB DD DSN=hlq.SEALINK, DISP=SHR
//SLSCNTL DD DSN=hlq.DBASEPRM, DISP=SHR
//SLSCNTL2 DD DSN=hlq.DBASESEC, DISP=SHR
//SLSSTBY DD DSN=hlq.DBASETBY, DISP=SHR
//SLSPRINT DD SYSOUT=*
//SLSIN DD *
CONFIG
GLOBAL MAXVTV=32000 MVCFREE=40
RECLAIM THRESHLD=70 MAXMVC=40 START=35
VTSS NAME=VSM501 LOW=70 HIGH=80 MAXMIG=8 RETAIN=5
RTD NAME=VSM52A00 DEVNO=2A00 CHANIF=0C
RTD NAME=VSM52A01 DEVNO=2A01 CHANIF=0D
RTD NAME=VSM52A02 DEVNO=2A02 CHANIF=0G
RTD NAME=VSM52A03 DEVNO=2A03 CHANIF=0H
RTD NAME=VSM52A04 DEVNO=2A04 CHANIF=0K
RTD NAME=VSM52A05 DEVNO=2A05 CHANIF=0L
RTD NAME=VSM52A06 DEVNO=2A06 CHANIF=0O
RTD NAME=VSM52A07 DEVNO=2A07 CHANIF=0P
RTD NAME=VSM52A08 DEVNO=2A08 CHANIF=1C
```

```

RTD  NAME=VSM52A09 DEVNO=2A09 CHANIF=1D
RTD  NAME=VSM52A0A DEVNO=2A0A CHANIF=1G
RTD  NAME=VSM52A0B DEVNO=2A0B CHANIF=1H
RTD  NAME=VSM52A0C DEVNO=2A0C CHANIF=1K
RTD  NAME=VSM52A0D DEVNO=2A0D CHANIF=1L
RTD  NAME=VSM52A0E DEVNO=2A0E CHANIF=1O
RTD  NAME=VSM52A0F DEVNO=2A0F CHANIF=1P
VTD  LOW=9900 HIGH=99FF


```

IOCP Example for Single MVS Host Connected to a VSM5 Using ESCON Directors

Figure I-4 shows a configuration diagram for a single MVS host connected to a VSM5 through ESCON Directors, and the example that follows it shows IOCP statements for this configuration. Note that:

- From MVSA, you define 8 CHPIDs, with each path switched in the ESCON Director, for a total of 8 channels running to the VSM5.
- You code 16 CNTLUNIT statements to define the VSM5 as 16 3490 images.
- You code IODEVICE statement to define the 16 VTDs that are associated with each 3490 image.

Figure I-4 Configuration Diagram: Single MVS Host Connected to a VSM5 Using ESCON Directors



```

ESCD4C CHPID PATH=(20,70),TYPE=CNC,SWITCH=4C
ESCD4D CHPID PATH=(21,71),TYPE=CNC,SWITCH=4D
ESCD4E CHPID PATH=(30,80),TYPE=CNC,SWITCH=4E
ESCD4F CHPID PATH=(31,81),TYPE=CNC,SWITCH=4F

```

```

CU1 CNTLUNIT CUNUMBR=001,
    PATH=(20,21,30,31,70,71,80,81),
    LINK=(D0,D4,D0,D4,D4,D0,D4,D0),
    UNIT=3490,CUADD=0,
    UNITADD=((00,16))

```

```

STRING1 IODEVICE ADDRESS=(0500,16),
        CUNUMBER=(001),
        UNIT=3490,
        UNITADD=00,STADET=Y

CU2 CNLUNIT CUNUMBR=002,
    PATH=(20,21,30,31,70,71,80,81),
    LINK=(D0,D4,D0,D4,D4,D0,D4,D0),
    UNIT=3490,CUADD=1,
    UNITADD=((00,16))

STRING2 IODEVICE ADDRESS=(0510,16),
        CUNUMBER=(002),
        UNIT=3490,
        UNITADD=00,STADET=Y
.
.
.
CU15 CNLUNIT CUNUMBR=015,
    PATH=(20,21,30,31,70,71,80,81),
    LINK=(D0,D4,D0,D4,D4,D0,D4,D0),
    UNIT=3490,CUADD=E,
    UNITADD=((00,16))

STRING15 IODEVICE ADDRESS=(05E0,16),
        CUNUMBER=(015),
        UNIT=3490,
        UNITADD=00,STADET=Y

CU16 CNLUNIT CUNUMBR=016,
    PATH=(20,21,30,31,70,71,80,81),
    LINK=(D0,D4,D0,D4,D4,D0,D4,D0),
    UNIT=3490,CUADD=F,
    UNITADD=((00,16))

STRING16 IODEVICE ADDRESS=(05F0,16),
        CUNUMBER=(016),
        UNIT=3490,
        UNITADD=00,STADET=Y

```

VSM5 New Models

The VSM5 provides new models that offer the capacities shown in [Table J-1](#).

Table J-1 VSM5 New Model Capacities

Drive Capacity	1-Array TBE	2-Arrays TBE	3-Arrays TBE	4-Arrays TBE	VSM Model	Product Family ID
450GB	NA	45	NA	NA	VSM5-45TB-IFF3	580
450GB	NA	NA	68	NA	VSM5-68TB-IFF3	580
450GB	NA	NA	NA	90	VSM5-90TB-IFF3	580
146GB	NA	7.5/11/14	NA	NA	VSM5-1.25TB-IFF3	567
146GB	NA	NA	16/18/21	NA	VSM5-16TB-IFF3	567
146GB	NA	NA	NA	23/25/28	VSM5-23TB-IFF3	567
146GB	.8/1.25	NA	NA	NA	VSM5E -.8TB	567

- TBE =The approximate maximum effective capacity in Terabytes (TB).
- Models VSM5-45TB-IFF3, VSM5-68TB-IFF3, and VSM5-90TB-IFF3 require PTF L1H14UN (SES7000).

These models can contain a maximum of 500,000 VTVs. VSM5 new models require the following VTSS microcode:

- VSM5 new models require the following VTSS microcode:
 - **For Models VSM5-45TB-IFF3, VSM5-68TB-IFF3, and VSM5-90TB-IFF3 only,** VTSS microcode level H02.07.
 - **For all other Models,** VTSS microcode level D02.07.
 - **Models VSM5-1.25TB-IFF3, VSM5-16TB-IFF3, and VSM5-23TB-IFF3** support ESCON channel cards. For more information, see "[VSM5 ESCON/FICON Configurations](#)."

[Table J-2](#) shows the supported channel card configurations for:

- VSM5 - All models including VSM5c but not including VSM5e and VSM5escon. For information on VSM5escon, see "[VSM5 ESCON/FICON Configurations](#)."
- VSM5e.

Table J-2 VSM5 New Model Supported Channel Card Configurations

VSM Model	Storage Cluster 0 2	Storage Cluster 0 3	Storage Cluster 0 4	Storage Cluster 0 5	Storage Cluster 1 12	Storage Cluster 1 13	Storage Cluster 1 14	Storage Cluster 1 15
5	VCF	NA	VCF	NA	VCF	NA	VCF	NA
5	VCF	VCF	VCF	NA	VCF	VCF	VCF	NA
5	VCF	VCF	VCF	VCF	VCF	VCF	VCF	VCF
5e	VCF	NA	VCF	NA	VCF	NA	VCF	NA
53	ICE	NA	ICE	NA	ICE	NA	ICE	NA

VSM 6/VSM 7 Configuration

VSM 6, supported by ELS 7.0 and above and VSM 7, supported by ELS 7.2 and above, provide greater connectivity, capacity, and throughput than VSM 5. [Table K-1](#) summarizes the ELS features.

Table K-1 ELS Features

Feature	Description
Host Interfaces	FICON
RTDs supported	RTD connections must be FICON channel.
VLE support	VLE TCP/IP connections are supported.
VTSS clustering support	TCP/IP CLINKs only are supported by any combination of VSM 5s, VSM 6s, and VSM 7s.
Maximum VTDs per VTSS	256 or optionally 512 for VSM7
Maximum VTVs per VTSS	4,000,000 (VSM6 and VSM7)

Reconfiguring VTCS

Note: You can now use the CDSCREAT utility to create a new CDS for Tapeless VSM.

To add ELSs to your configuration, you need to update the VTCS CONFIG file as follows:

1. Run the DECOMP utility to produce a CONFIG deck of your current configuration.

You update the CONFIG deck to add VSM 6s and VSM 7s, then submit the updated file to the CONFIG utility.

2. Update the CONFIG deck to add CONFIG VTSS statements for the VSM 6s and VSM 7s you are adding as shown in the following example.

```
VTSS NAME=VSM6A LOW=70 HIGH=80 MAXMIG=8 MINMIG=4 RETAIN=5
VTD LOW=8900 HIGH=89FF CMDCONN=TCPIP IPCONN=(192.80.12.1,192.80.12.2)
VTSS NAME=VSM6B LOW=70 HIGH=80 MAXMIG=8 MINMIG=4 RETAIN=5
VTD LOW=9900 HIGH=99FF CMDCONN=TCPIP IPCONN=(192.80.12.3,192.80.12.4)
```

In this example, the CMDCONN and IPCONN parameters define the VTSS to host connection as TCP/IP, so you do not have to do the tasks described in

"Configuration Examples for Single MVS Host Connected to a VSM 6/VSM 7 Using FICON Directors."

3. If necessary, add definitions for VTSS clustering.

You can only cluster VSM 6s/VSM 7s with other VSM 6s/VSM 7s or with VSM5s, and you must use TCP/IP connections for clustering (using the CONFIG CLINK IPIF parameter) as shown in the example of a cluster of two VSM 6s in the following example. Additionally, VSM 6 and VSM 7 can be part of a Cross-TapePlex Replication (CTR) configuration.

```
TAPEPLEX THISPLEX=TMSPA
VTSS NAME=VSM6A LOW=70 HIGH=80 MAXMIG=8 MINMIG=4 RETAIN=5
.
.
VTSS NAME=VSM6B LOW=70 HIGH=80 MAXMIG=8 MINMIG=4 RETAIN=5
.
.
CLUSTER NAME=CLUSTER6 VTSSs (VSM6A, VSM6B)
  CLINK VTSS=VSM6A IPIF=0I:0
  CLINK VTSS=VSM6A IPIF=1I:0
CLINK VTSS=VSM6B IPIF=0I:0
  CLINK VTSS=VSM6B IPIF=1I:1
CLINK VTSS=VSM6A IPIF=0A:0
  CLINK VTSS=VSM6A IPIF=1A:0
CLINK VTSS=VSM6B IPIF=0A:0
  CLINK VTSS=VSM6B IPIF=1A:0
```

Note: Valid values for CONFIG CLINK IPIF in the ci:p format are c=0 or 1, i=A or I, p=0 through 3 for VSM 5s, VSM 6s, and VSM 7s. For VSM5s, this value must match the value specified on the VSM5 IFF Configuration Status Screen. For VSM 6s and VSM 7s, this must be unique for each VTSS; and does not correspond to an actual value on the VSM 6/VSM 7 TCP/IP ports.

For more information on VTSS clustering, refer to the *ELS Disaster Recovery and Offsite Data Management Guide*.

If your VSM 6s/VSM 7s do not have RTDs attached, go to Step 5. Otherwise, continue with Step 4.

4. If necessary, add definitions for VSM 6/VSM 7-to-RTD connections.

VSM 6/VSM 7-to-RTD connections use FICON interfaces, and these connections are initially defined at the VSM 6/VSM 7.

Next, you define the paths to the RTDs using the CONFIG RTDPATH statement, as shown in the following example.

```
TAPEPLEX THISPLEX=TMSPA STORMNGR=VL1
VTSS NAME=VSM6A LOW=70 HIGH=80 MAXMIG=8 MINMIG=4 RETAIN=5
.
.
.
RTDPATH NAME=VSM6ARTD1 CHANIF=0A:0
RTDPATH NAME=VSM6ARTD2 CHANIF=0C:0
RTDPATH NAME=VSM6ARTD3 CHANIF=0E:0
RTDPATH NAME=VSM6ARTD4 CHANIF=0G:0
```


In this example, VSM6A has four FICON connections to RTDs.

Note: Valid values for CONFIG RTDPATH CHANIF in the ci:p format are c=0 or 1, i=A, C, E, G, I, K, M, or O, p=0 through 3 for VSM5s, VSM 6s and VSM 7s. For VSM5s, this value must match the actual FICON interface values. For VSM 6s and VSM 7s, this must be unique for each VTSS; and does not correspond to an actual value on the VSM 6/VSM 7 FICON ports.

If your VSM 6s/VSM 7s do not have VLEs attached, go to Step 7. Otherwise, continue with Step 5.

5. If necessary, add definitions for VSM 6/VSM 7-to-VLE connections.

For ELS 7.2 and later, VLE support is included in the base level. For ELS 7.0 and 7.1, get the latest SMP/E receive HOLDDATA and PTFs described in [Table K-2](#) and SMP/E APPLY with GROUPEXTEND.

Table K-2 ELS Supporting PTFs for VLE

ELS 7.0	ELS 7.1
L1H16C1	L1H16J6
L1H1672	L1H1674

For VTCS 7.0 and above, the CONFIG TAPEPLEX statement defines the TapePlex that VTCS is running under and provides the list of defined VLEs on the CONFIG TAPEPLEX STORMNGR parameter as shown in the following example.

```
TAPEPLEX THISPLEX=TMSPA STORMNGR=VL1
VTSS NAME=VTSS1 LOW=70 HIGH=80 MAXMIG=8 MINMIG=4 RETAIN=5
RTDPATH NAME=VL1RTD1 STORMNGR=VLE1 IPIF=0A:0
RTDPATH NAME=VL1RTD2 STORMNGR=VLE1 IPIF=0A:1
RTDPATH NAME=VL1RTD3 STORMNGR=VLE1 IPIF=0I:0
RTDPATH NAME=VL1RTD4 STORMNGR=VLE1 IPIF=0I:1
RTDPATH NAME=VL1RTD5 STORMNGR=VLE1 IPIF=1A:0
RTDPATH NAME=VL1RTD6 STORMNGR=VLE1 IPIF=1A:1
RTDPATH NAME=VL1RTD7 STORMNGR=VLE1 IPIF=1I:0
RTDPATH NAME=VL1RTD8 STORMNGR=VLE1 IPIF=1I:1
VTD LOW=6900 HIGH=69FF
```

In this example, note:

- The CONFIG TAPEPLEX statement, which defines TMSPA as the TapePlex that VTCS is running under and the names of all connected VLEs (which in this example is a single VLE called VLE1).
- The CONFIG RTDPATH statements, which define a single VLE RTD for each path from the VTSS to the VLE. In this example, the CONFIG RTDPATH statements for VTSS1 specify:
 - The name of the RTDPATH.
 - The connections to the defined VLEs (VLE1).
 - The IPIF value for each VTSS to VLE port connection in ci:p format where c is 0 or 1, i is A or I, and p is 0 through 3.

Note: For VSM5s, this value must match the values specified on the VSM5 IFF Configuration Status Screen. For VSM 6s and VSM 7s, this must be unique for each VTSS but does not correspond to an actual value on the VSM 6/VSM 7 TCP/IP ports.

6. If desired, define VLE virtual devices.

VTCS 7.1 and above systems can, of course, drive VLE as VTCS 7.0 does. In this mode, however, the number of VLE RTD targets is limited by the number of paths from a VTSS. Additionally, the VLE RTDs are assigned to fixed VTSS paths. The path from a VTSS to the VLE is always reserved by VTCS regardless of whether any VTSS to VLE data transfer is occurring.

However, with VTCS 7.1 and above (with the PTFs described in [Table K-2](#)), you can define a VLE with more VLE RTD targets than there are paths from the VTSS to the VLE, which means:

- The path from the VTSS to the VLE is not reserved unless a VTSS to VLE data transfer is required.
- More VLE RTD operations can occur simultaneously. For example, an audit of a VMVC does not require data transfer between the VTSS and the VLE.

As shown in the following example, the VLEs are defined using a CONFIG STORMNGR statement, not the CONFIG TAPEPLEX STORMNGR parameter. The CONFIG STORMNGR statement specifies the VLEs that VTCS connects to. Additionally, for each VLE, the CONFIG STORMNGR VLEDEV parameter defines the number and the names of the RTD devices that the VLE emulates. The more devices defined (up to the maximum of 96 devices per VLE), the greater the level of concurrent activities VTCS can schedule on the VLEs.

```
TAPEPLEX THISPLEX=TMSPC
STORMNGR NAME=VLE1 VLEDEV(S000-S05F)
STORMNGR NAME=VLE2 VLEDEV(S000-S05F)
VTSS NAME=VTSS1 LOW=70 HIGH=80 MAXMIG=8 MINMIG=4 RETAIN=5
RTDPATH NAME=VL1RTD1 STORMNGR=VLE1 IPIF=0A:0
RTDPATH NAME=VL1RTD2 STORMNGR=VLE1 IPIF=0A:1
RTDPATH NAME=VL1RTD3 STORMNGR=VLE1 IPIF=0I:0
RTDPATH NAME=VL1RTD4 STORMNGR=VLE1 IPIF=0I:1
RTDPATH NAME=VL1RTD5 STORMNGR=VLE2 IPIF=1A:0
RTDPATH NAME=VL1RTD6 STORMNGR=VLE2 IPIF=1A:1
RTDPATH NAME=VL1RTD7 STORMNGR=VLE2 IPIF=1I:0
RTDPATH NAME=VL1RTD8 STORMNGR=VLE2 IPIF=1I:1
VTD LOW=6900 HIGH=69FF
```

In this example, note:

- The CONFIG TAPEPLEX statement now simply defines TMSPC as the TapePlex that VTCS is running under. It does **not** define the connected VLEs.
- The CONFIG STORMNGR statements, which define the VLEs configured in this system - VLE1 and VLE2 in this example. These statements also specify the number of VLE devices using the VLEDEV parameter. In this example, each VLE has the maximum of 96 emulated devices, which allows VTCS to schedule up to 96 processes on each VLE. The VLE device addresses are in the form of Sxxx (where xxx is a hexadecimal value. For example, S000-S05F represents 96 emulated devices.
- The CONFIG RTDPATH statements for VTSS1, which specify:

- The name of the RTDPATH.
- The connections to the two ELSs.
- The IPIF value for each VTSS to ELS port connection in *ci:p* format where:
 - c* is 0 or 1.
 - i* is A or I.
 - p* is 0 through 3.

Note: For VSM5s, this value must match the values specified on the VSM5 IFF Configuration Status Screen. For VSM 6s and VSM 7s, this must be unique for each VTSS but does not correspond to an actual value on the VSM 6/VSM 7 TCP/IP ports.


7. If desired, you can use a VSM 6 and VSM 7 in a Cross-TapePlex Replication (CTR) configuration. For more information, refer to the *ELS Disaster Recovery and Offsite Data Management Guide*.
8. Finally, adjust your policies as needed. For more information, refer to *Configuring HSC and VTCS*.

Configuration Examples for Single MVS Host Connected to a VSM 6/VSM 7 Using FICON Directors

Figure K-1 shows a configuration for a single MVS host connected to a VSM 6 or VSM 7 using FICON Directors. For examples, see "[IOCP Example](#)" and "[Using HCD to Define a VSM 6/VSM 7](#)."

Note: If ESCON and FICON channels are configured to the same logical control unit, MVS issues message CBDG489I, which indicates that mixing ESCON and FICON channel paths on a logical control unit should be used only for the migration from ESCON to native FICON but should not be used permanently. This is a warning message only, and does not indicate an error.

Figure K-1 Configuration Diagram: Single MVS Host Connected to a VSM 6 Using FICON Directors



IOCP Example

Note the following:

- From MVSA, you define 8 CHPIDs, with each path switched in the FICON Director, for a total of 8 channels running to the VSM 6 or VSM 7.
- You code 16 CNTLUNIT statements to define the VSM 6/VSM 7 as 16 3490 images or if the VSM supports 512 VTDs, then code 32 CNTLUNIT statements. The 17th CNTLUNIT CUADD number would be 10 and the 32nd CNTLUNIT number would be 1F.
- You code IODEVICE statements to define 16 VTDs that are associated with each 3490 image or if the VSM supports 512 VTDs, then code 32 IODEVICE statements.

```
FCD4C CHPID PATH=(20,70),TYPE=FC,SWITCH=4C
FCD4D CHPID PATH=(21,71),TYPE=FC,SWITCH=4D
FCD4E CHPID PATH=(30,80),TYPE=FC,SWITCH=4E
FCD4F CHPID PATH=(31,81),TYPE=FC,SWITCH=4F
```

```
CU1 CNTLUNIT CUNUMBR=001,
    PATH=(20,21,30,31,70,71,80,81),
    LINK=(D0,D4,D0,D4,D4,D0,D4,D0),
    UNIT=3490,CUADD=0,
    UNITADD=((00,16))
```

```
STRING1 IODEVICE ADDRESS=(0500,16),
    CUNUMBR=(001),
    UNIT=3490,
    UNITADD=00,STADET=Y
```

```
CU2 CNTLUNIT CUNUMBR=002,
    PATH=(20,21,30,31,70,71,80,81),
    LINK=(D0,D4,D0,D4,D4,D0,D4,D0),
    UNIT=3490,CUADD=1,
    UNITADD=((00,16))
```

```

STRING2 IODEVICE ADDRESS=(0510,16),
        CUNUMBER=(002),
        UNITADD=00,STADET=YUNIT=3490,
        .
        .
        .

CU15 CNLUNIT CUNUMBR=015,
      PATH=(20,21,30,31,70,71,80,81),
      LINK=(D0,D4,D0,D4,D4,D0,D4,D0),
      UNIT=3490,CUADD=E,
      UNITADD=((00,16))

STRING15 IODEVICE ADDRESS=(05E0,16),
         CUNUMBER=(015),
         UNIT=3490,
         UNITADD=00,STADET=Y

CU16 CNLUNIT CUNUMBR=016,
      PATH=(20,21,30,31,70,71,80,81),
      LINK=(D0,D4,D0,D4,D4,D0,D4,D0),
      UNIT=3490,CUADD=F,
      UNITADD=((00,16))

STRING16 IODEVICE ADDRESS=(05F0,16),
         CUNUMBER=(016),
         UNIT=3490,
         UNITADD=00,STADET=Y
    
```

Note: Unlike ESCON, FICON supports multiple active I/Os per channel. If the number of active VTDs is less than the number of channels configured to the VTSS, the I/Os to those VTDs may not be evenly spread across all the channels. As the number of active VTDs increases to be greater than the number of channels configured to the VTSS, the channel subsystem will spread the I/Os across all the channels. If it is desired to spread the I/Os across all of the channels even when only a few VTDs are active, it is necessary to use the preferred path feature to force the channel subsystem to spread the I/Os across the channels. The preferred path feature is specified using the PATH= parameter on the IODEVICE statement. When you specify preferred path on the IODEVICE statement, the channel subsystem always tries the preferred path first. If it is busy or unavailable, the channel subsystem next tries the channel path following the preferred path in the rotation order, and so on.

The following example shows IODEVICE statements for STRING1 without using preferred pathing.

```

STRING1 IODEVICE ADDRESS=(0500,16),
        CUNUMBER=(001),
        UNIT=3490,
        UNITADD=00,STADET=Y
    
```

The following example shows IODEVICE statements for STRING1 using preferred pathing. If you are using preferred pathing, you need to use these kind of IODEVICE statements for all paths.

```
STRING10 IODEVICE ADDRESS=(0500,2),  
          CUNUMBER=(001),  
          UNIT=3490,  
          UNITADD=00,STADDET=Y,  
          PATH=20
```

```
STRING12 IODEVICE ADDRESS=(0502,2),  
          CUNUMBER=(001),  
          UNIT=3490,  
          UNITADD=00,STADDET=Y,  
          PATH=21
```

```
STRING14 IODEVICE ADDRESS=(0504,2),  
          CUNUMBER=(001),  
          UNIT=3490,  
          UNITADD=00,STADDET=Y,  
          PATH=30
```

```
STRING16 IODEVICE ADDRESS=(0506,2),  
          CUNUMBER=(001),  
          UNIT=3490,  
          UNITADD=00,STADDET=Y,  
          PATH=31
```

```
STRING18 IODEVICE ADDRESS=(0508,2),  
          CUNUMBER=(001),  
          UNIT=3490,  
          UNITADD=00,STADDET=Y,  
          PATH=70
```

```
STRING1A IODEVICE ADDRESS=(050A,2),  
          CUNUMBER=(001),  
          UNIT=3490,  
          UNITADD=00,STADDET=Y,  
          PATH=71
```

```
STRING1C IODEVICE ADDRESS=(050C,2),  
          CUNUMBER=(001),  
          UNIT=3490,  
          UNITADD=00,STADDET=Y,  
          PATH=80
```

```
STRING1E IODEVICE ADDRESS=(050E,2),  
          CUNUMBER=(001),  
          UNIT=3490,  
          UNITADD=00,STADDET=Y,  
          PATH=81
```

Using HCD to Define a VSM 6/VSM 7

The following sections tell how to use HCD to define a VSM 6 or VSM 7.

- ["Define the FICON Channels"](#)
- ["Define the VSM 6/VSM 7 As Sixteen 3490 Images and Connect Them to the FICON Channels"](#)

- "Define 16 VTDs for Each Control Unit"

Note:

- The following sections show procedures for using HCD to define a VSM 6 or VSM 7, not the actual HCD panels. For more information, refer to the IBM HCD documentation.
 - You may not need to do all of the above tasks (if, for example, you have already defined the FICON channels).
-
-

Define the FICON Channels

Using the values shown in [Figure K-1](#), define (and dedicate) the FICON channels to MVSA, with each path switched in the FICON Director, for a total of 8 channels running to the VSM 6 or VSM 7 as follows:

1. Select a processor from the Processor List panel.
2. Select the Work with attached channel paths action from the context menu.
3. Press PF11 to display the Add Channel Path panel.
4. Use the Add Channel Path panel to add 8 channels to the processor.

For each channel, complete the channel path details as shown in [Table K-3](#). Assuming that the FICON switches are pre-defined, enter the dynamic switch number and the entry switch number and port number. The entry port is the port where the channel is physically plugged in. The dynamic switch is the switch number of the FICON Director where the dynamic connection takes place. In the case of single switches in each channel path, the entry switch ID is the same as the dynamic switch ID.

Table K-3 Add Channel Path Values

Configuration Mode	Number of CHPIDs	Channel Path Type	Operation Mode	Channel Path ID	Dynamic Switch ID	Entry Switch ID	Entry Port
LPAR	1	FC	DED	20	4C	4C	F0
LPAR	1	FC	DED	70	4C	4C	F1
LPAR	1	FC	DED	21	4D	4D	F0
LPAR	1	FC	DED	71	4D	4D	F1
LPAR	1	FC	DED	30	4E	4E	F0
LPAR	1	FC	DED	80	4E	4E	F1
LPAR	1	FC	DED	31	4F	4F	F0
LPAR	1	FC	DED	81	4F	4F	F1

5. For each channel, the Define Access List panel is displayed.
Select MVSA from the Partition Name Number Usage Description list.

Define the VSM 6/VSM 7 As Sixteen 3490 Images and Connect Them to the FICON Channels

Define the 3490 images and connect them to the FICON channels as follows:

1. Select Control Units from the Define, Modify or View Configuration Data panel.

2. On the Control Unit List panel, Press PF11 to display the Add Control Unit panel.
For each control unit, complete the control unit details as shown in [Table K-4](#).

Note: HCD may provide default values for the protocol and I/O concurrency level fields, depending on the control unit type being defined.

VSM7 VTSS support 512 VTDs with current maintenance. Control unit numbers 011 – 020 and CUADD number 10 – 1F are for VTDs 256 – 511. If your VSM does not support 512 VTDs, these Control Units do not need to be coded.

Table K-4 Add Control Unit Values

Control Unit Type	Channel Path IDs	Control Unit Number	Unit Address	Number of Units	Logical Address = CUADD
3490	20, 70, 21, 71, 30, 80, 31, 81	001	00	16	0
3490	20, 70, 21, 71, 30, 80, 31, 81	002	00	16	1
3490	20, 70, 21, 71, 30, 80, 31, 81	003	00	16	2
3490	20, 70, 21, 71, 30, 80, 31, 81	004	00	16	3
3490	20, 70, 21, 71, 30, 80, 31, 81	005	00	16	4
3490	20, 70, 21, 71, 30, 80, 31, 81	006	00	16	5
3490	20, 70, 21, 71, 30, 80, 31, 81	007	00	16	6
3490	20, 70, 21, 71, 30, 80, 31, 81	008	00	16	7
3490	20, 70, 21, 71, 30, 80, 31, 81	009	00	16	8
3490	20, 70, 21, 71, 30, 80, 31, 81	00A	00	16	9
3490	20, 70, 21, 71, 30, 80, 31, 81	00B	00	16	A
3490	20, 70, 21, 71, 30, 80, 31, 81	00C	00	16	B

Table K-4 (Cont.) Add Control Unit Values

Control Unit Type	Channel Path IDs	Control Unit Number	Unit Address	Number of Units	Logical Address = CUADD
3490	20, 70, 21, 71, 30, 80, 31, 81	00D	00	16	C
3490	20, 70, 21, 71, 30, 80, 31, 81	00E	00	16	D
3490	20, 70, 21, 71, 30, 80, 31, 81	00F	00	16	E
3490	20, 70, 21, 71, 30, 80, 31, 81	010	00	16	F
3490	20, 70, 21, 71, 30, 80, 31, 81	011	00	16	10
3490	20, 70, 21, 71, 30, 80, 31, 81	012	00	16	11
3490	20, 70, 21, 71, 30, 80, 31, 81	013	00	16	12
3490	20, 70, 21, 71, 30, 80, 31, 81	014	00	16	13
3490	20, 70, 21, 71, 30, 80, 31, 81	015	00	16	14
3490	20, 70, 21, 71, 30, 80, 31, 81	016	00	16	15
3490	20, 70, 21, 71, 30, 80, 31, 81	017	00	16	16
3490	20, 70, 21, 71, 30, 80, 31, 81	018	00	16	17
3490	20, 70, 21, 71, 30, 80, 31, 81	019	00	16	18
3490	20, 70, 21, 71, 30, 80, 31, 81	01A	00	16	19
3490	20, 70, 21, 71, 30, 80, 31, 81	01B	00	16	1A
3490	20, 70, 21, 71, 30, 80, 31, 81	01C	00	16	1B

Table K-4 (Cont.) Add Control Unit Values

Control Unit Type	Channel Path IDs	Control Unit Number	Unit Address	Number of Units	Logical Address = CUADD
3490	20, 70, 21, 71, 30, 80, 31, 81	01D	00	16	1C
3490	20, 70, 21, 71, 30, 80, 31, 81	01E	00	16	1D
3490	20, 70, 21, 71, 30, 80, 31, 81	01F	00	16	1E
3490	20, 70, 21, 71, 30, 80, 31, 81	020	00	16	1F

When you are finished, press Enter to display the Select Processor / Control Unit panel.

3. Select MVSA from the Select Processor / Control Unit panel and the Select (connect, change) action on the context menu.

When you are finished, press Enter to display the Select Processor / Control Unit panel. Note that Yes in the Att column indicates that the control unit is connected to the processor in the same row.

Define 16 VTDs for Each Control Unit

Define 16 VTDs and connect them to each control unit as follows:

1. On the Define, Modify or View Configuration data panel, select option I/O devices to display the I/O Device List panel.
2. Press F11 to display the Add Devices panel.

Complete the Add Devices panel. When you add a range of devices and you specify the first device serial number or a general description, you are prompted for the serial numbers of the other devices and can modify the description of each device.

VSM7 VTSS support 512 VTDs with current maintenance. Device numbers 0600 – 06F0 and CUADD numbers 10 – 1F are for VTDs 256 – 511. If your VSM does not support 512 VTDs, these devices do not need to be coded.

Table K-5 Add Devices Values

Device Type	Device Number	Unit Address	Number of Units	Logical Address = CUADD
3390	0500	00	16	0
3390	0510	00	16	1
3390	0520	00	16	2
3390	0530	00	16	3
3390	0540	00	16	4

Table K-5 (Cont.) Add Devices Values

Device Type	Device Number	Unit Address	Number of Units	Logical Address = CUADD
3390	0550	00	16	5
3390	0560	00	16	6
3390	0570	00	16	7
3390	0580	00	16	8
3390	0590	00	16	9
3390	05A0	00	16	A
3390	05B0	00	16	B
3390	05C0	00	16	C
3390	05D0	00	16	D
3390	05E0	00	16	E
3390	05F0	00	16	F
3390	0600	00	16	10
3390	0610	00	16	11
3390	0620	00	16	12
3390	0630	00	16	13
3390	0640	00	16	14
3390	0650	00	16	15
3390	0660	00	16	16
3390	0670	00	16	17
3390	0680	00	16	18
3390	0690	00	16	19
3390	06A0	00	16	1A
3390	06B0	00	16	1B
3390	06C0	00	16	1C
3390	06D0	00	16	1D
3390	06E0	00	16	1E
3390	06F0	00	16	1F

- Pressing Enter to display the Define Device to Operating System Configuration panel.

This panel lets you specify the operating systems to which the devices are to be defined. Select the Config ID, and the Select (connect, change) action to connect the devices to each operating system.

Note that YES is later shown in the Defined column to indicate that this device has been defined for the operating system in the same row.

Pressing Enter displays the Define Device Parameters / Features panel.

This panel allows device features to be specified to the selected operating system. Review the following MVS operating system features:

- The OFFLINE feature defaults to the value specified in the UIM (which may not be appropriate for tape devices and ESCON directors).
- The DYNAMIC feature defaults to Yes for devices that allow dynamic reconfiguration. Check your subsystems and applications that manage their own devices to make sure they support dynamically modified device control blocks (UCBs).

The status of a device may be changed from DYNAMIC=YES to DYNAMIC=NO, and vice versa, by a software-only dynamic change to a configuration. When changing from DYNAMIC=NO to DYNAMIC=YES, this must be the only change made to the device at that time. You are allowed, however, to make changes to a device that is currently defined as DYNAMIC=YES and at the same time change it to DYNAMIC=NO.

Note: The DYNAMIC parameter is shown only when the appropriate device supports the dynamic I/O configuration function.

- The SHARED feature defaults to No. Since many MVS installations have more than one image, this parameter should be carefully checked and changed to "Yes" if appropriate.
4. Pressing Enter displays the Assign/Unassign Device to Esoteric panel.
 5. Specify Yes or No to assign or unassign devices to the esoteric.
If you do not want to assign all the devices currently being defined to this esoteric, you can limit the devices being assigned by specifying a starting device number and the number of devices to be assigned.
 6. Press Enter to go back through the panels to the Device List panel.



Tapeless VSM

"Tapeless VSM" basically means that you can have a VTSS without any RTDs directly attached to the VTSS; in the CONFIG deck, there are no RTD statements for the tapeless VTSS. Tapeless VSM applies to VSM4s and VSM5s.

How Does Tapeless VSM Work?

Configuring and managing a Tapeless VSM works as follows:

1. In the CONFIG deck, there are no RTD statements for the Tapeless VTSS.

Note: For clustered VTSS configurations, all VTSSs in the cluster must be Tapeless or all VTSSs in the cluster must have RTDs attached. You cannot mix Tapeless VTSSs and VTSSs with RTDs attached within a cluster.

2. The new MGMTCLAS NOMIGRAT parameter specifies that VTVs in the Management Class are not candidates for migration, consolidation or export, but are candidates to reside on a tapeless VTSS.

VTSS selection is changed to prefer Tapeless VTSSs for VTVs in Management Classes with NOMIGRAT, and to disallow VTVs without NOMIGRAT from VTSSs with no RTDs.

NOMIGRAT parameter is mutually exclusive with the ACLIST, IMMEDMIG, DUPLEX, MIGPOL, ARCHAGE, ARCHPOL, RESTIME, CONSRC and CONTGT parameters.

3. A Management Class can specify DELSCR(YES), which is a proactive method to cause VSM to delete scratched VTVs, which frees VTSS buffer space and (logically) deletes any VTV copies from MVCs so that MVC space can be reclaimed. As an alternative, you can specify DELSCR(NO) and use the DELETSCR utility (which now provides a VTSS parameter to scratch VTVs on a per VTSS basis) to do demand deletes of scratched VTVs.

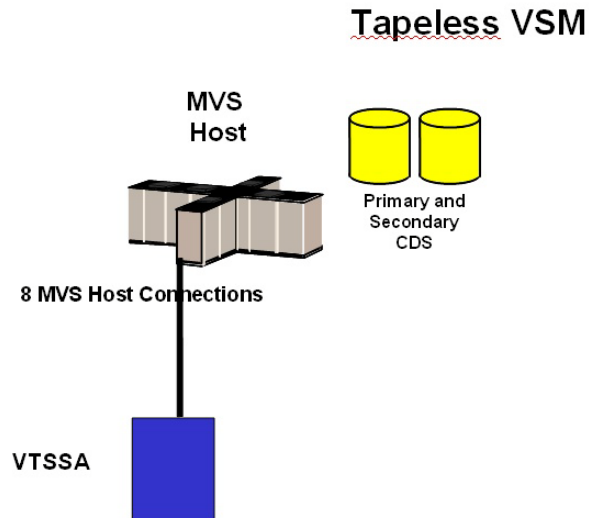
For a sample configuration, see ["Tapeless VSM Example."](#)

Note: If you have an environment that is completely Tapeless (no RTDs attached to any VTSS system), then in your LIBGEN you need to code a dummy ACS as shown in the example in ["LIBGEN Example for Tapeless ACS."](#)

Tapeless VSM Example

Figure L-1 shows a simple Tapeless VSM configuration where VTSSA has no RTDs attached.

Figure L-1 Tapeless VSM



Configuring the System

To configure the example system shown in Figure L-1, code a CONFIG deck as shown in the following example.

In this example, note that there are no RTD statements for VTSSA because it is Tapeless.

```
//CREATECF EXEC PGM=SWSADMIN, PARM='MIXED'
//STEPLIB DD DSN=hlq.SLSLINK, DISP=SHR
//SLSCNTL DD DSN=FEDB.VSMLMULT.DBASEPRM, DISP=SHR
//SLSCNTL2 DD DSN=FEDB.VSMLMULT.DBASESEC, DISP=SHR
//SLSSTBY DD DSN=FEDB.VSMLMULT.DBASETBY, DISP=SHR
//SLSPRINT DD SYSOUT=*
//SLSIN DD *
CONFIG RESET CDSLEVEL(V62ABOVE)
GLOBAL MAXVTV=32000 MVCFREE=40 VTVATTR=SCRATCH RECALWER=YES LOCKSTR=VTCS_LOCKS
REPLICAT=CHANGED VTVPAGE=LARGE MAXRTDS=32
RECLAIM THRESHLD=70 MAXMVC=40 START=35
VTVVOL LOW=905000 HIGH=999999 SCRATCH
VTVVOL LOW=C00000 HIGH=C25000 SCRATCH
VTVVOL LOW=RMM000 HIGH=RMM020 SCRATCH
MVCVOL LOW=N25980 HIGH=N25989
MVCVOL LOW=N35000 HIGH=N35999
```

```
VTSS NAME=VTSSA
VTD LOW=8900 HIGH=89FF
```

Defining Policies

To define policies for the example system shown in [Figure L-2](#), do the following:

1. Enable the Advanced Management Feature.
2. Create a Storage Class for VTSSA.

```
STOR NAME(TAPEL)
```

3. Create the Management Classes that point to the Storage Classes in Step 2.

```
MGMT NAME(TAPEL) NOMIGRAT DELSCR(YES)
```

In this example you created a Management Class that points to the corresponding Storage Class created in Step 2.

Note: In addition to NOMIGRAT, the Management Class in this example specifies DELSCR(YES), which is a proactive method to cause VSM to delete scratched VTVs, which frees VTSS buffer space and (logically) deletes any VTV copies from MVCs so that MVC space can be reclaimed. As an alternative, you can specify DELSCR(NO) and use the DELETSCR utility (which now provides a VTSS parameter to scratch VTVs on a per VTSS basis) to do demand deletes of scratched VTVs.

4. Create an SMC Policy that specifies virtual media and assigns the Management Class created in Step 3.

```
POLICY NAME(PTAPEL) MEDIA(VIRTUAL) MGMT(TAPEL)
```

5. Create a TAPEREQ statement to route data to VSM and assign the "Tapeless" Policy to the data.

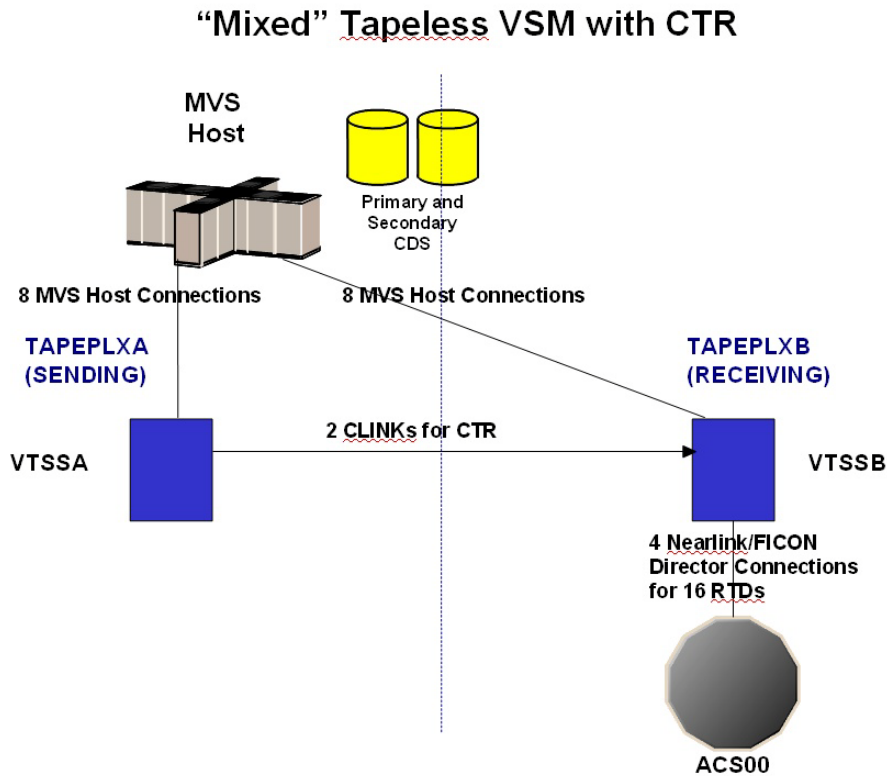
In this example, the TAPEREQ statement specifies to route data sets with HLQ mask *.HR.** to VSM and assign Policy PTAPEL.

Note: Although you can use SMC policies to direct your migrations to a specific esoteric, StorageTek recommends using only MGMTCLAS so that the SMC/VTCS allocation influencing can use any VTSS that supports the MGMTCLAS requirements.

Using CTR in “Tapeless” and “With Tape” Environments

[Figure L-2](#) shows two non-clustered VTSSs (VTSSA and VTSSB in separate TapePlexes) where VTSSA has no RTDs attached. VTSSA sends VTVs to VTSSB using CTR so that VTSSB can migrate the VTVs to tape in ACS00.

Figure L-2 Mixed Tapeless VSM



As shown in Figure L-2, therefore, you can construct a "Tapeless" TapePlex, use CTR to connect this TapePlex to another TapePlex that has RTDs attached. With this configuration, you can make a "TapeLess" copy in the first TapePlex and replicate it to the second TapePlex where it can be migrated to tape. For more information on DTR, refer to the *ELS Disaster Recovery and Offsite Data Management Guide*.

LIBGEN Example for Tapeless ACS

The following examples show a LIBGEN example for a Tapeless ACS, where the SLISTATN ADDRESS=(0032) statement denotes a dummy ACS. Note that for new Tapeless VSM configurations, you can alternatively use the CDSCREAT utility as described in "Running the SLUADMIN CDSCREAT Utility to Create the CDS (Tapeless VSM Only)."

```

*
* LABELS WILL BE GENERATED IN THE OUTPUT LIBGEN AS FOLLOWS:
* ACS      - "ACSXX"   WHERE "XX" IS THE HEX ACS NUMBER 00-FF
*           STARTING WITH ZERO
* LSM      - "LSMXXYY" WHERE "XX" IS THE HEX ACS NUMBER OF THIS LSM
*           AND "YY" IS THE HEX LSM NUMBER (00-FF) IN THAT
*           ACS, STARTING AT ZERO FOR EACH NEW ACS
* STATION  - "STXXH"  WHERE "XX" IS THE HEX ACS NUMBER AND H IS THE
*           HOST INDEX IN HEX (0-F)
* PANEL    - "PXXYYPP" WHERE "XX" IS THE HEX ACS NUMBER, YY IS THE HEX
*           LSM NUMBER, AND PP IS THE DECIMAL PANEL NUMBER
*           OF THE DRIVE PANEL
* DRIVE    - "DXXYYPPH" WHERE "XX" IS THE HEX ACS NUMBER, YY IS THE HEX
*           LSM NUMBER, PP IS THE DRIVE PANEL NUMBER IN DEC
*

```



```

*                               IMAL, AND H IS THE HOST INDEX IN HEX
*
LIBGEN  SLIRCVRY TCHNIQ=NONE
*
      SLILIBRY SMF=231,                      X
          ACSLIST=ACSLIST,                   X
          HOSTID=(EC20,EC21),                 X
          MAJNAME=STKALSQN,                   X
          CLNPRFX=CLN,                         X
          COMPRFX=!,                           X
          DRVHOST=,                             X
          SCRLABL=SL
*
ACSLIST  SLIALIST ACS00
*
ACS00    SLIACS LSM=(LSM0000,LSM0001,LSM0002,LSM0003),  X
          STATION=(ST000,ST001)
*
ST000    SLISTATN ADDRESS=(0032)
ST001    SLISTATN ADDRESS=(0032)
*
LSM0000  SLILSM PASTHRU=((0,M),(0,M),(0,M)),           X
          ADJACNT=(LSM0001,LSM0002,LSM0003),           X
          DRIVE=(1),                                     X
          DRVELST=(P000001),                             X
          TYPE=8500,                                     X
          DOOR=8500-2
*
P000001  SLIDLIST HOSTDRV=(D0000010,D0000010)
*
D0000010 SLIDRIVS ADDRESS=(,,,,,,,,,,,,,)
*
LSM0001  SLILSM PASTHRU=((0,S),(0,M),(0,M)),           X
          ADJACNT=(LSM0000,LSM0002,LSM0003),           X
          DRIVE=(1),                                     X
          DRVELST=(P000101),                             X
          TYPE=8500,                                     X
          DOOR=8500-2
*
P000101  SLIDLIST HOSTDRV=(D0001010,D0001010)
*
D0001010 SLIDRIVS ADDRESS=(,,,,,,,,,,,,,)
*
LSM0002  SLILSM PASTHRU=((0,S),(0,S),(0,M)),           X
          ADJACNT=(LSM0000,LSM0001,LSM0003),           X
          DRIVE=(1),                                     X
          DRVELST=(P000201),                             X
          TYPE=8500,                                     X
          DOOR=8500-2
*
P000201  SLIDLIST HOSTDRV=(D0002010,D0002010)
*
D0002010 SLIDRIVS ADDRESS=(,,,,,,,,,,,,,)
*
LSM0003  SLILSM PASTHRU=((0,S),(0,S),(0,S)),           X
          ADJACNT=(LSM0001,LSM0002,LSM0000),           X
          DRIVE=(1),                                     X
          DRVELST=(P000301),                             X
          TYPE=8500,                                     X
          DOOR=8500-2

```

```
*  
P000301 SLIDLIST HOSTDRV=(D0003010,D0003010)  
*  
D0003010 SLIDRIVS ADDRESS=(,,,,,,,,,,,,,  
*  
          SLIENDGN ,
```

Running the oVTCS CDS Database Server

The oVTCS CDS database server component proxy enables a client oVTCS running on the VSM console server to act as a local z/OS VTCS so that it can access the z/OS-resident CDS database.

The oVTCS client must have its own "host slot" within the CDS to send and receive broadcast messages, hold locks, and so forth. This requirement imposes the following limitations on executing the oVTCS CDS database server:

- The oVTCS CDS database server cannot execute on a z/OS host that has VTCS executing or that may have VTCS executing. The HSC subsystem must be started with the NOVTCS startup parameter specified in the EXEC statement.
- The oVTCS CDS database server is restricted to communicating with a single oVTCS client; there is a 1-to-1 relationship between client and server. For example, if there are two oVTCS appliances, they each require their own HSC host. As a result, there are two instances of HSC each executing their own oVTCS CDS database server as shown in [Figure M-1](#):

Figure M-1 Relationship Between oVTCS Client and oVTCS CDS DB Server



The "Server Host ID" boxes in [Figure M-1](#) represent the server on which the DBSERVer command is issued (see "DBSERVer command" in the *ELS Command, Control Statement, and Utility Reference*). It is also the host that supplies the proxy host ID for the oVTCS client CDS. Thus, oVTCS1 is host ID MVSA and oVTCS2 is host ID MVSB. Note that MVSA and MVSB may share the same CDS, but they must have separate host IDs.

- The oVTCS CDS database server requires HSC database services be active. Therefore, the oVTCS CDS database server cannot be started before the HSC has reached BASE service level.
- You can include the DBSERVer START command in the HSC startup parameter file, or you can issue it as an HSC operator command. It is not a UUI/XAPI enabled command.
- The oVTCS CDS database server requires its own TCP/IP port assignment for its socket listener. The port number is specified on the DBSERVer START command. If

you execute the SMC HTTP server on the same host as the oVTCS CDS database server, you must specify different port numbers.

- If requested by Oracle StorageTek Software support, use GTF and the TRace VTcs command to enable tracing of the oVTCS CDS database server on z/OS.

Glossary

Note: Glossary entries that include the (I) designation are credited to the *IBM Dictionary of Computing*.

4410

Oracle's StorageTek standard Library Storage Module (LSM).

4480

Oracle's StorageTek 18-track 1/2-inch cartridge transport.

4490

Oracle's StorageTek 36-track long-tape cartridge transport with ESCON support. Also known as Silverton.

9310

Oracle's StorageTek Library Storage Module (LSM), a high-performance version of the standard 4410 LSM. Also known as PowderHorn.

9360

Oracle's StorageTek Library Storage Module (LSM). Also known as WolfCreek.

9740

Oracle's StorageTek Library Storage Module (LSM). Also known as TimberWolf.

Abnormal end of task (abend)

A software or hardware problem that terminates a computer processing task.

ACSid

A method used in the LIBGEN process to identify ACSs by using decimal digits, 00 to 99.

ACS

Automated Cartridge System, a fully-automated, cartridge storage and retrieval library subsystem consisting of one or more Library Storage Modules (LSMs) connected by pass-thru ports.

ACS library

A library is composed of one or more Automated Cartridge Systems (ACSs), attached cartridge drives, and cartridges residing in the ACSs.

ACSLs

Automated Cartridge System Library Software, Oracle's StorageTek library control software, which runs in the UNIX®-based Library Control System.

address

Coded representation of hardware id, or the destination or origination of data.

allocation

The assignment of resources to a specific task.

asynchronous transmission

Character-oriented data transmission (as distinct from IBM's block-mode transmission).

Automated Cartridge System (ACS)

See [ACS](#).

Automated Cartridge System Library Software (ACSLs)

See [ACSLs](#).

automatic mode

A relationship between an LSM and all attached hosts. LSMs operating in automatic mode handle cartridges without operator intervention. This is the normal operating mode of an LSM that has been modified online. The opposite situation is "manual mode." See manual mode.

bar code

A code consisting of a series of bars of varying widths. This code appears on the external label attached to the spine of a cartridge and is equivalent to the volume serial number (volser). This code is read by the robot's machine vision system.

BISYNC

Binary Synchronous Communications. An early low-level protocol developed by IBM and used to transmit data on a synchronous communications link. It is a form of data transmission in which synchronization of characters is controlled by timing signals generated at the sending and receiving stations.

CAPid

A CAPid uniquely defines the location of a CAP by the LSM on which it resides. A CAPid is of the form "AA:LL:CC" where AA:LL is the LSMid and CC is the two digit CAP number.

cartridge

The plastic housing around the tape. It is approximately 4 inches (100 mm) by 5 inches (125 mm) by 1 inch (25 mm). The tape is threaded automatically when loaded in a transport. A plastic leader block is attached to the tape for automatic threading. The spine of the cartridge contains an OCR/Bar Code label listing the VOLSER (tape volume identifier).

Cartridge Access Port (CAP)

An assembly that allows several cartridges to be inserted into or ejected from an LSM without human entry into the LSM.

cartridge drive (CD)

A hardware device containing two or four cartridge transports and associated power and pneumatic supplies.

cartridge tape I/O driver

Operating system software that issues commands (for example, read, write, and rewind) to cartridge subsystems. It is the software focal point for attaching a particular type of control unit. (An example is Oracle's StorageTek CARTLIB product.)

cell

A receptacle in the LSM in which a single cartridge is stored.

CGI

Common Gateway Interface

channel

A device that connects the host and main storage with the input and output devices' control units. A full-duplex channel has two paths (that is, 2 wires, or one wire with signals at two frequencies). A half-duplex channel requires that one port receives while the other transmits.

channel-to-channel (CTC)

Refers to the communication (transfer of data) between programs on opposite sides of a channel-to-channel adapter.(I)

client

The ultimate user of the ACS services as provided by the Library Control System.

client link

The communications link between the LCS and a client.

client/server

A model of interaction in a distributed system in which a program at one site serves a request to a program at another site and awaits a response. The requesting program is called a client; the program satisfying the request is called a server.

client system

The system to which the LCS provides an interface to a StorageTek Automated Cartridge System.

Client System Component (CSC)

Software that provides an interface between the Client Computing System's operating system and the StorageTek Library Control System (LCS).

coaxial cable

A transmission medium used in data transmissions for networks using synchronous communications, as opposed to twisted-pair, the primary medium for asynchronous RS-232 communications.

complex

A system composed of other systems, specifically the ACS server system and the client system.

connected mode

A relationship between a host and an ACS. In this mode, the host and an ACS are capable of communicating (in the sense that at least one station to this ACS is online).

connection number

The unique identifier on the server for a communications path. The number is assigned by TCP/IP to identify the unique connection between the server node and a specific port on the server, and the client node and a specific port on the client. The connection number exists only as long as the connection exists.

console

The primary I/O device to control a session on a system.

control data set (CDS)

The data set used by the host software to control the functions of the automated library. Also called a library database.

Control Path Adaptor (CPA)

A Bus-Tech, Inc. hardware device that allows communications between a host processor's block multiplexer channel and a local area network.

Control Unit (CU)

A microprocessor-based unit situated locally between a channel and an I/O device. It translates channel commands into device commands and sends device status to the channel.

coupling facility

A special logical partition that provides high-speed caching, list processing, and locking functions in a sysplex.(I)

coupling facility channel

A high bandwidth fiber optic channel that provides the high-speed connectivity required for data sharing between a coupling facility and the central processor complexes directly attached to it.(I)

CTC

Channel-to-channel.

Data Path Adapter

A hardware device which translates from a client computing system's data protocol to the data protocol of the StorageTek Control Unit or IMU. An example is DEC's TC44-AA/BA STI-to-4400 ACS Interconnect.

data set

A set of records treated as a unit.

data sharing

The ability of concurrent subsystems or application programs to directly access and change the same data while maintaining data integrity.(I)

device number

A four-digit hexadecimal number that uniquely identifies a device attached to a processor.

device preferencing

The process of preferring one 36-track transport type over another 36-track transport type.

device separation

See [drive exclusion](#).

DFSMS

Data Facility Storage Management Subsystem.

direct access storage device (DASD)

IBM's term for a disk drive storage device.

directed allocation

See drive prioritization.

disconnected mode

A relationship between a host and an ACS. In this mode, the host and the ACS are not capable of communicating (there are no online stations to this ACS).

dotted-decimal notation

The syntactic representation of a 32-bit integer that consists of four 8-bit numbers written in base ten with periods (dots) separating them. In TCP/IP descriptions, dotted-decimal notation is used for Internet addresses.

drive exclusion

(previously known as device separation) refers to the SMC function of excluding drives for an allocation request based on SMC exclusion criteria.

drive panel

An LSM wall containing tape transports. The drive panel for a T9840 transport can contain either 10 or 20 transports. The drive panel for a non-T9840 transport can contain a maximum of 4 transports.

drive prioritization

(previously known as directed allocation) refers to the SMC function of influencing selection of a particular drive based on allocation criteria, including volume location.

Dual LMU

A hardware/microcode feature that provides a redundant LMU capability.

dump

A printed representation of the contents of main storage at time t. This representation is used for debugging purposes.

dynamic server switching

The capability of switching server processors when a system failure occurs on the active server.

ECART

Enhanced Capacity Cartridge.

Enhanced Capacity Cartridge

A cartridge that has a length of 1100 feet and can be used only on 36-track transports (that is, 4490, 9490, and 9490EE).

Enterprise Systems Connection (ESCON)

A set of products and services that provides a dynamically connected environment using optical cables as a transmission medium.(I)

error codes (EC)

Numeric codes displayed by messages indicating the type of problem that caused an error.

error recovery procedures (ERP)

Procedures designed to help isolate and, where possible, to recover from errors in equipment.

ESCON

Enterprise Systems Connection.

esoteric name

The name assigned to transports that have the same device type.

Ethernet

One LAN architecture using a bus topology that allows a variety of computers to be connected to a common shielded coaxial spine. The Ethernet architecture is similar to the IEEE 802.3 standard.

event control block (ECB)

Provides an area for a completion code to be stored when an operation has completed.

file

A set of related records treated as a unit.

File Transfer Protocol (FTP)

A TCP/IP command that provides a way to transfer files between machines connected through TCP/IP.

foreign socket

One of two end-points in a TCP/IP connection-oriented protocol. Specifies the address of a foreign host that can connect to the server.

GB

1,073,741,834 bytes of storage

handshake

A flow-of-control signal sent by one process to another.

helical cartridge

A high capacity, helical scan cartridge that can hold up to 50GB of uncompressed data. This cartridge can be used only on RedWood (SD-3) transports.

host computer

A computer that controls a network of computers.

HTTP

Hypertext Transfer Protocol.

IEEE 802.3

A standard produced by the IEEE and accepted worldwide for local area networks using CSMA/CD (Carrier Sense Multiple Access with Collision Detection).

ICRC

Improved Cartridge Recording Capacity. A compression and compaction feature that increases the amount of data that can be stored on a 1/2-inch cartridge.

initial program load (IPL)

A process that activates a machine reset.

Interactive Storage Management Facility

A series of applications for defining DFSMS/MVS storage groups and classes.

Internet

A collection of networks using TCP/IP that functions as a virtual network.

Internat address

The numbering system used to specify a network or host on that network for TCP/IP communications. Standard Internet address notation is dotted-decimal format.

Internet Protocol (IP)

Formal description of messages and rules two networks use to exchange messages.

ISMF

Interactive Storage Management Facility

job control language (JCL)

A problem oriented language designed to describe a job's processing requirements to an operating system.

JES

Job entry subsystem.(I)

JES2

An MVS subsystem that receives jobs into the system, converts them to internal format, selects them for execution, processes their output, and purges them from the system. In an installation with more than one processor, each JES2 processor independently controls its job input, scheduling, and output processing. See also JES3.(I)

JES3

An MVS subsystem that receives jobs into the system, converts them to internal format, selects them for execution, processes their output, and purges them from the system. In complexes that have several loosely coupled processing units, the JES3 program manages processors so that the global processor exercises centralized control over the local processors and distributes jobs to them through a common job queue. See also JES2.(I)

LAN

Local Area Network, a network in a small (local) geographic area

LCS

Library Control System (LCS)

LCS processor console

The Library Control System processor console is used to control the VM operating system (for the VM-based LCS).

LCU

See Library Control Unit.

LIBGEN

The process of defining the configuration of a library to the VM/HSC.

library

See TapePlex.

library cartridge transport

See transport.

library complex

A library complex consists of one HSC Control Data Set (CDS) and may contain up to 256 Automatic Cartridge Systems (ACSs), each of which may contain up to 24 Library Storage Modules (LSMs).

library control component

Software that controls the mounting and dismounting of cartridges in an ACS.

library control platform

The hardware and software that provides the proper environment for the Library Control System.

library control processor

Properly configured computer hardware that supports the operation of the Library Control System.

Library Control Software

A library control component, the client system interface, and library utilities.

Library Control System (LCS)

The library control platform and the Library Control Software.

Library Control Unit (LCU)

The portion of an LSM that controls the movements of the robot.

library database

A file or data set containing information about the location and status of the removable media volumes, such as cell location, scratch status. Also called a control data set (CDS).

library drive

A cartridge drive in the ACS, as distinct from a standalone cartridge drive.

Library Management Unit (LMU)

See [LMU](#).

library mode

The operation of a 4480 Cartridge Subsystem as part of a 4400 Automated Cartridge System, as opposed to manual mode, in which the operator inserts cartridges into the transports. See manual mode.

Library Storage Module (LSM)

See [LSM](#).

LMU

Library Management Unit, A hardware and software product that coordinates the activities of one or more LSMs/LCUs.

local area network (LAN)

See [LAN](#).

local port

The designation of a given application or process among many that are available for a TCP/IP-capable host processor.

local socket

The address combination of a TCP/IP-capable host's network address and a specific port for an application process.

logical port (LP)

See [LP](#).

LP

logical port, CLS software that interfaces with the client system. The CLSLP is one of the software components used to pass data between the client system and the VM/HSC.

LSM

Library Storage Module, The standard LSM (4410) a twelve-sided structure with storage space for up to around 6000 cartridges. It also contains a free-standing, vision-assisted robot that moves the cartridges between their storage cells and attached transports. See also PowderHorn, SL3000, SL8500, and WolfCreek.

LSMid

An LSMid is composed of the ACSid joined to (concatenated with) the LSM number.

LSM number

A method used to identify an LSM. An LSM number is the result of defining the SLIACS macro LSM parameter during a LIBGEN. The first LSM listed in this parameter acquires the LSM number of 00 (decimal) the second LSM listed acquires a number of 01, and so forth, until all LSMs are identified (up to a maximum of 99).

manual mode

Operation of a cartridge drive apart from an ACS. See library mode.

master LMU

The LMU currently controlling the functional work of the ACS in a dual LMU configuration.

mixed configuration

A configuration that contains different types of cartridge drives in both manual and library modes.

modem

A device that enables digital data to be transmitted over an analog transmission facility.

multi-client

The environment where more than one (homogenous or heterogeneous) client system is connected to one LCS.

MVS system console

The MVS/CSC provides an operator interface through the MVS system console.

OCR label

Optical character recognition label. An external label attached to the spine of a cartridge that is both human and machine readable.

operator console

In this document, the operator console refers to the MVS client system console.

operating system (OS)

Software that controls the execution of programs that facilitate overall system operation.

Pass-thru Port (PTP)

A mechanism that allows a cartridge to be passed from one LSM to another in a multiple LSM ACS.

physical port

The communications hardware required to support a server/client link.

physical volume

A physically bound unit of data file media. See cartridge.

PowderHorn (9310)

The high-performance version of the standard LSM.

preconfigured package

A storage server package including all hardware, software, and configuration parameter settings delivered by the vendor.

product change request (PCR)

A request for enhancement to a product. Normally, this request comes from a client, but may come from Oracle.

program temporary fix (PTF)

A software release designed to remedy one or a series of defects.

program update tape (PUT)

One or more tapes containing updates to, or new versions of, the MVS/CSC system software.

protocol

A formal description of message formats and the rules two or more machines must follow to exchange these messages.

recovery

Automatic or manual procedures to resolve problems in the server system.

reel-id

Identifier of a specific tape volume. Equivalent to volume serial number (VOLSER).

request

Term used to refer to commands issued to the 4400 ACS to perform a tape-related function.

scratch tape

A tape that is available to any user because it is not owned.

scratch tape subpool

A defined subset of all scratch tapes. Subpools are composed of one or more ranges of volsers with similar physical characteristics (type of volume—reel or cartridge, reel size, length, physical location, and so on). Some installations may also subdivide their scratch pools by other characteristics such as label type.

SD-3

Oracle's StorageTek helical cartridge transport. Also known as RedWood.

SL3000

Oracle's StorageTek SL3000 modular library offers mixed media, logical and physical partitioning capabilities, advanced management, and high availability. It supports mixed environments, including mainframe and open systems, and is scalable from 200 to just under 6,000 cartridge slots.

SL8500

Oracle's StorageTek SL8500 modular library offers mixed media, logical and physical partitioning capabilities, advanced management, and high capacity and availability. It supports mixed environments, including mainframe and open systems, and is scalable from the standard 1,450 cartridge slots to 100,880 cartridge slots in a complex configuration.

server

An ELS library control system such as HSC. In SMC a server is represented by a named SERVER path to a named TAPEPLEX. While the SMC HTTP server component is required as the middle ware on the remote host, the server, as far as SMC is concerned is the ELS library control system operating on the remote host.

socket

A unique address on a network plus a node address plus the id of one specific application on a specific network. An abstraction used by TCP/IP.

standard capacity cartridge

A cartridge that can be used on any longitudinal transport (that is, 4480, 4490, 9490, or 9490EE).

standby

The status of a station that has been varied online but is connected to the standby LMU of a dual LMU ACS.

standby LMU

The redundant LMU in a dual LMU configuration that is ready to take over in a Master LMU failure or when the operator issues a SWITCH command.

station

A hardware path between the host computer and an LMU over which the VM/HSC and LMU send control information.

storage server

A set of hardware and software products designed to enable heterogeneous computer systems to use automated tape cartridge library services.

Storage Management Component (SMC)

Software interface between IBM's z/OS operating system and StorageTek automated library control systems, HSC and MVS/CSC. SMC performs the allocation processing, message handling, and SMS processing for the ELS solution.

switchover

The assumption of master LMU function by the standby LMU.

synchronous

See [BISYNC](#).

synchronous LAN

Local area network built on synchronous communications.

Systems Network Architecture (SNA)

A description of the logical structure, formats, protocols, and operational sequences for transmitting information units through and controlling the configuration and operation of networks.

T9840A

Oracle's StorageTek access-centric cartridge transport capable of reading and writing 9840A cartridges.

T9840B

Oracle's StorageTek access-centric cartridge transport capable of reading and writing T9840B cartridges.

T9840C

Oracle's StorageTek access-centric cartridge transport capable of reading and writing T9840C cartridges.

T9840D

Oracle's StorageTek access-centric cartridge transport capable of reading and writing T9840D cartridges.

T9940A

Oracle's StorageTek capacity-centric cartridge transport capable of reading and writing 60GB T9940A cartridges.

T9940B

Oracle's StorageTek capacity-centric cartridge transport capable of reading and writing 200GB T9940B cartridges.

T10000A

Oracle's StorageTek T10000 A high-capacity cartridge transport capable of reading and writing 120GB or 500GB T10000A cartridges.

T10000B

Oracle's StorageTek T10000 B high-capacity cartridge transport capable of reading and writing 240GB or 1TB T10000B cartridges.

T10000C

Oracle's StorageTek T10000 C high speed/capacity tape drive, delivering up to 252 MB/sec and 5 TB, native, making it ideal for data center operations with growing data volumes.

T10000D

Oracle's StorageTek T10000D high speed/capacity tape drive, delivering up to 252 MB/sec and 8.5 TB native capacity, making it ideal for data center operations with growing data retention requirements.

tape drive

A tape processing device consisting of up to four transports in a cabinet. A drive can refer to an individual transport.

tape library management system (TLMS)

TLMS, as used in this document, refers to any tape library management system, not to CA-1.

TapePlex

(formerly "library"), a single StorageTek hardware configuration, normally represented by a single HSC Control Data Set (CDS). A TapePlex may contain multiple Automated Cartridge Systems (ACSs) and Virtual Tape Storage Subsystems (VTSSs).

TCP/IP

Transmission Control Protocol/Internet Protocol.

trace event type

Types of event traced through the system when tracing is enabled.

trace file

A file that contains information useful for debugging the system.

transaction

A specific set of input that triggers the execution of a specific process.

Transmission Control Protocol (TCP)

An inter-network standard protocol that provides a full-duplex stream service.

transport

An electro-mechanical device used to thread, position, and read or write from a tape.

UCB

Unit Control Block.

userid

Sometimes referred to as the VM userid, the userid is the name that identifies a specific "virtual machine" user or client.

utility

Program that performs a function ancillary to the chief function(s) of a computer system.

virtual machine (VM)

See [VM](#).

virtual storage

A feature of the OS where main storage requirements are allocated by segments (or pages) as needed by programs, thus creating the apparent existence of unlimited or virtual storage.

Virtual Storage Manager (VSM)

A storage solution that virtualizes volumes and transports in a VTSS buffer to improve media and transport use.

Virtual Tape Control System (VTCS)

The primary host code for the Virtual Storage Manager (VSM) solution. This code operates in a separate address space, but communicates closely with HSC.

Virtual Tape Storage Subsystem (VTSS)

The DASD buffer containing virtual volumes (VTVs) and virtual drives (VTDs). The VTSS is a StorageTek RAID 6 hardware device with microcode that enables transport emulation. The RAID device can read and write "tape" data from/to disk, and can read and write the data from/to a real tape drive (RTD).

Virtual Telecommunications Access Method (VTAM)

IBM host-resident communications software that serves as a common interface for communications.

VM

Virtual Machine, A functional simulation of a computer and its associated devices. Each virtual machine is controlled by a suitable operating system.

VM/SP or VM/XA

A proprietary operating system of IBM corporation that consists mainly of two major components, CP and CMS.

volume

A tape cartridge (data carrier) that is mounted or dismounted as a unit.

volume serial number (VOLSER)

An identifier of a physical volume.

WolfCreek (9360)

The high-performance LSM with a smaller capacity than the standard LSM.

ZCART

An extended-enhanced cartridge that uses a thinner media to provide twice the capacity of the enhanced capacity (ECART) cartridge. This cartridge has a length of 2200 feet and can be used only on TimberLine 9490EE 36-track transports.

B

backing out from an I-level CDS, 1-18

Bulk CAP

adding to or removing from a partition with the HSC down, 5-54

adding to or removing from a partition with the HSC up, 5-55

changing the CAP state for Bulk CAP with the HSC up, 5-55

changing the CAP state with the HSC down, 5-54

C

CDS

oVTCS database server, M-1

configuring HSC

creating and cataloging the HSC startup procedure, 4-19

creating the HSC CDS, 4-1

creating the HSC PARMLIB member, 4-17

defining command authority, 4-19

defining the SYS1.PARMLIB member

SMFPRMxx, 4-15, 4-16

defining volumes to ELS, 4-10

performance considerations, 4-24

running multiple HSC tapeplexes on one

LPAR, 4-10

updating HSM, 4-19

configuring VTCS

building a simple CONFIG deck, 6-2

defining MVC pool volser authority, 6-5

storing VTCS locks in a coupling facility, 6-6

updating the TMS for VSM, 6-4

control data set

limiting access privileges, 1-6

control data set (CDS)

backing out from an I-level CDS, 1-18

placement, 1-4

serializing, 1-6

sharing, 1-4

space requirements, 1-14

types, 1-3

using CDS logging to recover the CDS, 1-18

VTCS level requirements, 1-16

D

data chaining a VTD Read forward or Write command, 1-31

determining HSC/VTCS configuration values, 1-3

G

guidelines for MVCs, 1-23

H

HSC

starting, 7-1

stopping, 7-5

L

limiting CDS access privileges, 1-6

LSMs

planning for, 1-19

M

MVC media preferencing

using the STORclas MEDIA parameter for, 1-26

MVCs

guidelines for, 1-23

O

oVTCS CDS database server

running, M-1

P

partitioning

adding or removing a Bulk CAP to or from a partition

HSC down, 5-54

HSC up, 5-55

CAP considerations, 5-22

LibraryStation considerations, 5-22

requirements and prerequisites, 5-21

restrictions, 5-21

Streamline libraries, 5-20

term definitions, 5-23

- partitioning procedures
 - adding a library to the HSC complex, 5-33
 - adding a partition to the library, 5-38
 - adding resources to a library, 5-47
 - changing the CAP state for Bulk CAP with the HSC down, 5-54
 - changing the CAP state for Bulk CAP with the HSC up, 5-55
 - converting a non-partitioned library to a partitioned library, 5-28
 - converting a partitioned library to a non-partitioned library, 5-30
 - moving a partition from one host group to another host group, 5-44
 - removing a library from the HSC complex, 5-36
 - removing a partition from the library, 5-41
 - removing resources from a library, 5-51
 - starting from scratch using LIBGEN, SLICR, 5-26
 - varying the ACS online after a configuration change, 5-55
- planning for HSC/VTCS configuration
 - determining configuration values, 1-3
 - how VSM measures sizes and capacities, 1-2
- planning for library volumes, 1-23
- planning for library-attached transports, 1-21
- planning for LSMs, 1-19
- planning for VTDs, 1-29
- planning for VTSSs, 1-28
- planning for VTVs, 1-30
- planning VTCS operating policies
 - general policies, 2-10
 - locked resource policy, 2-2
 - making CDS logging required, 2-1
 - MVC policies, 2-29
 - VTSS policies, 2-2
 - VTV migration policies, 2-14
 - VTV replication policies, 2-25
- preparing for configuration
 - configuring MSP device numbers and esoterics, 3-1
 - defining a security system user ID for HSC/VTCS, 3-1
 - setting the MSP missing interrupt handler (MIH) value, 3-2
- procedures, dynamic hardware reconfiguration
 - adding/removing expansion modules, 5-16, 5-17
 - changing CAP types, HSC down, 5-18

R

- reconfiguring a tapeplex
 - adding a new VTSS, 5-20
 - combining two separate ACSs into one ACS, 5-7
 - combining two separate Streamline libraries, 5-8
 - consolidating two data centers, 5-6
 - HSC SET CDKLOCK utility, 5-20
 - HSC SET VAULTVOL and SET VAULT utilities, 5-20
 - partitioning Streamline libraries, 5-20
 - removing PTPs and splitting an ACS using

- MERGEcds, 5-11
- removing vaulted volume information, 5-12
- renaming a VTSS, 5-11
- using Audit instead of MERGEcds for Streamline libraries, 5-19
- using dynamic reconfiguration, 5-12
- using the MERGEcds utility, 5-2
- when do you need to reconfigure, 5-1
- without using MERGEcds, 5-12
- running the oVTCS CDS database server, M-1

S

- SL8500 library
 - changing CAP types
 - HSC down, 5-18
 - Near Continuous Operation (NCO)
 - adding/removing expansion modules, 5-16, 5-17
- SL8500 partitioning
 - adding or removing a Bulk CAP to or from a partition
 - HSC down, 5-54
 - HSC up, 5-55
- space requirements for a CDS, 1-14
- starting HSC, 7-1
- stopping HSC, 7-5

U

- using CDS logging to recover the CDS, 1-18

V

- VSM
 - configuration
 - record, M-1
 - how it measures sizes and capacities, 1-2
- VSM console
 - running the oVTCS CDS database server, M-1
- VTCS CDS level requirements, 1-16